

The digitalization of energy systems: towards higher energy efficiency

Présentée le 29 juillet 2021

Faculté des sciences et techniques de l'ingénieur
Laboratoire de photovoltaïque et couches minces électroniques
Programme doctoral en énergie

pour l'obtention du grade de Docteur ès Sciences

par

Marina DOROKHOVA

Acceptée sur proposition du jury

Prof. J. A. Schiffmann, président du jury
Prof. C. Ballif, Dr N. Würsch, directeurs de thèse
Prof. L. Wehenkel, rapporteur
Dr R. Rudel, rapporteur
Dr S.-R. Cherkaoui, rapporteur

Abstract

The energy industry and the electric power sector, in particular, are going through challenging times of disruptive changes, finding themselves at the crossroads of decarbonization, decentralization, and digitalization. Enduring extensive transformations in the midst of combating climate change, the energy value chain is restructuring itself to accommodate the growing penetration of renewables, increasing number of independent power producers, and augmented self-consumption. Therefore, new energy management approaches are required to accomplish energy transition, empower efficient interactions between the energy industry's key stakeholders, and ensure stable, reliable, and secure operations of the electric grid.

With the Fourth Industrial Revolution underway, it becomes evident that digitalization, driven by the explosion of data, improvements in computing power and capacity, widespread coverage of communication networks, and progress in machine learning, is the only way to increase energy efficiency in a sustainable and affordable manner. However, hindered by the energy industry's massive inertia, most energy utilities are missing the real momentum for unleashing large-scale digitalization and taking bold actions to embrace the Information and Communication Technology (ICT). Therefore, by proposing a set of ICT-supported software applications, models, and tools that operate on top of existing physical infrastructure, this thesis aims to bridge the gap between strategic roadmaps focused on the energy industry's digitalization and their actual implementation in real-world scenarios through digital energy services. Specifically, we conduct the research within the designed ICT-based smart building and smart community frameworks, the modular structure and scalability of which can serve as a backbone for future implementations of various digital energy management solutions.

The introduction of a novel unsupervised load disaggregation approach helps to raise the awareness of one's energy-related behavior and understand what drives the energy usage in residential households without compromising privacy and security. Showcasing algorithm's performance on real-world datasets from Norway and Germany highlights compliance with state-of-the-art disaggregation accuracy and reduced computational costs. The development of machine learning-based supervised and unsupervised building occupancy forecasting algorithms with prediction accuracies beyond 97% helps to identify best-suited windows for energy-saving opportunities and deliver insights on one's presence and absence patterns. Built on top of that occupancy-centric rule-based heating and air conditioning automation algorithm strives to unlock the buildings' massive potential for energy savings without compromising the occupants' thermal comfort. Simulations on real-world datasets collected in Portugal demonstrate the potential for more than 15% energy savings.

Zooming out from smart buildings towards smart communities, we focus on the important role of intelligent green mobility in supporting further digitalization of the electric power sector. To overcome the inconveniences posed by the sparsity of charging infrastructure and facilitate the adoption of Electric Vehicles (EVs), we present a reinforcement learning (RL)-based EV-specific routing method that guarantees paths' energy feasibility in a graph-theoretical context. Consequently, we propose several deep RL algorithms to control EV charging with the aim to increase renewables' self-consumption and EV drivers' satisfaction. Benchmarking against rule-based and model predictive control demonstrates RL's superior computational performance and better fitness for future mobility systems. Finally, we introduce an innovative decentralized blockchain-supported framework that enables secure and reliable accounting of energy exchanges within the smart community. Implementing it in a demonstration site in Switzerland shows blockchain's potential to reduce EV charging costs, transform the market's business model, and facilitate the large-scale deployment of EVs. This thesis concludes with recommendations to key energy stakeholders on further advancing the industry's digitalization.

Keywords: artificial intelligence, blockchain, digitalization, electric vehicle, energy efficiency, energy management system, information and communication technology, machine learning, photovoltaic, reinforcement learning, smart building, smart community

Résumé

L'industrie de l'énergie et le secteur de l'électricité, en particulier, traversent une période de changements perturbateurs, induits par les effets de la décarbonisation, de la décentralisation et de la numérisation. Dans le cadre de la lutte contre le changement climatique, la chaîne de valeur énergétique se restructure pour s'adapter à la pénétration croissante des énergies renouvelables, au nombre croissant de producteurs d'électricité indépendants et à l'augmentation de l'autoconsommation. Par conséquent, de nouvelles approches de gestion de l'énergie sont nécessaires pour réaliser la transition énergétique, favoriser des interactions efficaces entre les principaux acteurs de l'industrie énergétique et garantir un fonctionnement stable, fiable et sûr du réseau électrique.

Avec la quatrième révolution industrielle en cours, la numérisation est stimulée par l'explosion des données, les améliorations de la puissance et de la capacité de calcul, la couverture étendue des réseaux de communication et les progrès de l'apprentissage automatique. Il devient évident que cette numérisation est le seul moyen d'accroître l'efficacité énergétique de manière durable et abordable. Cependant, la plupart des services publics de l'énergie sont freinés par l'inertie massive du secteur de l'énergie et passent ainsi à côté du véritable élan qui permettrait de déclencher une numérisation à grande échelle et de prendre des mesures audacieuses pour adopter les technologies de l'information et de la communication (ICT). Par conséquent, cette thèse vise à combler le fossé entre les feuilles de route stratégiques axées sur la numérisation de l'industrie de l'énergie et leur mise en œuvre effective dans des scénarios du monde réel par le biais de services énergétiques numériques. Dans ce but, elle propose un ensemble d'applications, de modèles et d'outils logiciels supportés par les ICT qui fonctionnent au-dessus de l'infrastructure physique existante. Plus précisément, elle s'inscrit dans le cadre de bâtiments intelligents et de communautés intelligentes basés sur les ICT, dont la structure modulaire et l'évolutivité peuvent servir d'épine dorsale pour les futures mises en œuvre de diverses solutions de gestion numérique de l'énergie.

L'introduction d'une nouvelle approche non supervisée de désagrégation de la charge permet de sensibiliser les gens à leur comportement en matière d'énergie et de comprendre ce qui détermine la consommation d'énergie dans les foyers résidentiels sans compromettre la confidentialité et la sécurité. La présentation des performances de l'algorithme sur des ensembles de données réelles provenant de Norvège et d'Allemagne souligne la conformité de la précision de désagrégation et la réduction des coûts de calcul par rapport à l'état-de-l'art. Le développement d'algorithmes de prévision d'occupation de bâtiments, supervisés et non supervisés, basés sur l'apprentissage automatique, a permis d'atteindre une précision de

prédiction supérieure à 97%. Ces algorithmes permettent d'identifier les fenêtres les mieux adaptées aux possibilités d'économie d'énergie et de fournir des informations sur les habitudes de présence et d'absence des personnes. L'algorithme d'automatisation du chauffage et de la climatisation, basé sur des règles centrées sur l'occupation, s'efforce de libérer l'énorme potentiel d'économies d'énergie des bâtiments sans compromettre le confort thermique des occupants. Des simulations sur des ensembles de données réelles recueillies au Portugal démontrent un potentiel d'économies d'énergie de plus de 15%.

En passant des bâtiments intelligents aux communautés intelligentes, nous nous concentrons sur le rôle important de la mobilité verte intelligente pour soutenir la numérisation du secteur de l'énergie électrique. Pour surmonter les inconvénients posés par la rareté des infrastructures de recharge et faciliter l'adoption des véhicules électriques (EVs), nous présentons une méthode de routage spécifique aux EV basée sur l'apprentissage par renforcement (RL) qui garantit la faisabilité énergétique des chemins dans un contexte de théorie des graphes. En conséquence, nous proposons plusieurs algorithmes d'apprentissage par renforcement pour contrôler la recharge des EV dans le but d'augmenter l'autoconsommation des énergies renouvelables et la satisfaction des conducteurs de EV. La comparaison par rapport au contrôle prédictif basé sur des règles et des modèles démontre la performance de calcul supérieure du RL et sa meilleure adéquation aux futurs systèmes de mobilité. Enfin, nous présentons un cadre novateur décentralisé soutenu par la *blockchain*, qui permet une comptabilité sûre et fiable des échanges d'énergie au sein de la communauté intelligente. Sa mise en œuvre sur un site de démonstration en Suisse montre le potentiel de la *blockchain* pour réduire les coûts de recharge des EVs, transformer le modèle économique du marché et faciliter le déploiement à grande échelle des EVs. Cette thèse se conclut par des recommandations aux principaux acteurs du secteur de l'énergie pour faire progresser la numérisation de l'industrie.

Mots-clés : intelligence artificielle, *blockchain*, numérisation, véhicule électrique, efficacité énergétique, système de gestion de l'énergie, technologies de l'information et de la communication, apprentissage automatique, photovoltaïque, apprentissage par renforcement, bâtiment intelligent, communauté intelligente

Аннотация

Энергетика и, в частности, электроэнергетика претерпевает кардинальные изменения, обусловленные процессами декарбонизации, децентрализации и цифровизации. В условиях масштабных преобразований, направленных на борьбу с изменением климата, цепочка создания стоимости в энергетике трансформируется под влиянием ускоренного развития возобновляемых источников энергии, увеличения числа независимых производителей электроэнергии и повышения уровня самопотребления. Все эти факторы требуют новых подходов к управлению энергией, способных расширить возможности эффективного взаимодействия между ключевыми игроками энергетической отрасли и обеспечить стабильную, надёжную и безопасную эксплуатацию электросетей.

В контексте четвёртой промышленной революции становится очевидным, что цифровизация, активно развивающаяся благодаря значительному росту объёма данных, увеличению вычислительных скоростей и мощностей, широкому внедрению телекоммуникационных технологий и прогрессу в машинном обучении, является единственным способом доступного и устойчивого повышения энергоэффективности. Однако из-за огромной инерции энергетической отрасли большинству энергетических компаний не хватает действенного импульса для развёртывания крупномасштабной цифровизации и принятия смелых решений по внедрению информационных и коммуникационных технологий (ИКТ). Чтобы преодолеть существующий разрыв между стратегиями цифровизации в энергетике и их подлинной реализацией, данная диссертация предлагает набор приложений, моделей и инструментов, поддерживаемых ИКТ, которые могут применяться в качестве цифровых энергетических услуг. Представленные научные исследования проведены в рамках систем интеллектуального здания и интеллектуального сообщества, разработанных на базе ИКТ. Предложенная модульная структура и масштабируемость таких систем может служить прочным фундаментом для внедрения различных методов цифрового управления энергопотреблением.

В частности, предложение нового подхода к деагрегации электрической нагрузки, основанного на алгоритмах машинного обучения без учителя, помогает повысить осведомлённость людей об их поведении, связанном с потреблением энергии, и понять основные факторы, влияющие на энергопотребление в жилых домах, без ущерба для конфиденциальности и безопасности. Демонстрация производительности алгоритма на данных полученных из Норвегии и Германии подчёркивает соответствие

предложенного метода современной точности дезагрегирования, а также снижение вычислительных затрат. Кроме того, в диссертации описаны разработанные на основе машинного обучения с учителем и без учителя алгоритмы прогнозирования присутствия людей в зданиях. Такие алгоритмы, с 97% точностью предсказаний, помогают выявить благоприятные периоды для реализации механизмов снижения энергопотребления, без ущерба комфорту посетителей. Созданный на их основе алгоритм автоматизации отопления и кондиционирования позволяет поддерживать комфортные температурные условия в помещениях и раскрыть огромный потенциал зданий по уменьшению энергозатрат. Анализ применимости разработанного алгоритма, проведённый на данных из Португалии, демонстрирует потенциал экономии электроэнергии более чем 15%.

Переходя от умных зданий к умным сообществам, мы фокусируемся на важной роли интеллектуальных электрических видов транспорта в поддержке дальнейшей цифровизации электроэнергетического сектора. Чтобы преодолеть барьеры, связанные с редкостью зарядной инфраструктуры для электромобилей, и ускорить повсеместное внедрение экологического транспорта, в диссертации предлагается новый метод решения проблем дорожной навигации. Разработанный специально для электромобилей на основе машинного обучения с подкреплением, данный метод позволяет гарантировать энергетическую осуществимость маршрутов в дорожных сетях, созданных с помощью теории графов. Чтобы повысить самопотребление электроэнергии, полученной из возобновляемых источников, а также сделать использование электромобилей более комфортным, предлагаются несколько алгоритмов управления процессами зарядки электрического транспорта. Основанные на методах глубокого машинного обучения с подкреплением, созданные алгоритмы, в сравнении с другими популярными методами управления, демонстрируют повышенную вычислительную производительность и лучшую пригодность для будущих транспортных систем. В заключительной части исследований представлена инновационная децентрализованная система на основе технологии блокчейн, которая обеспечивает безопасный и надёжный учёт обмена энергией в интеллектуальном сообществе. Пилотный запуск данной системы в Швейцарии демонстрирует высокий потенциал блокчейна для снижения затрат на зарядку электромобилей, преобразования бизнес-модели рынка и содействия распространению электрического транспорта. Завершается диссертация рекомендациями для ключевых заинтересованных сторон в сфере энергетики по дальнейшему развитию цифровизации в отрасли.

Ключевые слова: искусственный интеллект, блокчейн, цифровизация, электромобиль, энергетическая эффективность, система управления энергией, информационные и коммуникационные технологии, машинное обучение, фотовольтаика, обучение с подкреплением, умное здание, умное сообщество

Abstract

L'industria energetica e il settore dell'energia elettrica, in particolare, stanno attraversando difficili cambiamenti dovuti alla decarbonizzazione, della decentralizzazione e della digitalizzazione. Nel pieno della lotta contro il cambiamento climatico, la *value chain* dell'energia si sta ristrutturando per accogliere la crescente penetrazione delle energie rinnovabili, il numero crescente di produttori di energia indipendenti e l'aumento dell'autoconsumo. Pertanto, sono necessari nuovi approcci di gestione dell'energia per realizzare la transizione energetica, potenziare le interazioni efficienti tra gli attori chiave dell'industria energetica e garantire operazioni stabili, affidabili e sicure della rete elettrica.

Con la quarta rivoluzione industriale in corso, diventa evidente che la digitalizzazione, guidata dall'esplosione dei dati, dai miglioramenti della potenza e della capacità di calcolo, dalla copertura diffusa delle reti di comunicazione e dai progressi nel *machine learning*, è l'unico modo per aumentare l'efficienza energetica in modo sostenibile e conveniente. Tuttavia, ostacolati dalla massiccia inerzia dell'industria energetica, la maggior parte dei servizi energetici stanno perdendo l'opportunità di sviluppare la digitalizzazione su larga scala e intraprendere azioni coraggiose per abbracciare le Tecnologie dell'Informazione e della Comunicazione (TIC). Pertanto, proponendo una serie di applicazioni software, modelli e strumenti supportati dal TIC e che operano in cima alle infrastrutture fisiche esistenti, questa tesi mira a colmare il divario tra le strategie focalizzate sulla digitalizzazione dell'industria energetica e la loro effettiva implementazione in scenari reali attraverso servizi energetici digitali. La ricerca è condotta all'interno di appositi framework basati sul TIC e sviluppati per edifici e comunità intelligenti. La struttura modulare dei framework e la loro scalabilità può servire come fondamenta per le future implementazioni di varie soluzioni di gestione digitale dell'energia.

L'introduzione di un nuovo approccio non supervisionato per la disaggregazione del carico, aiuta a capire cosa guida l'uso dell'energia nelle case residenziali e ad aumentare la consapevolezza di un uso sostenibile dell'energia senza compromettere la privacy e la sicurezza. Mostrando le prestazioni dell'algoritmo su dati provenienti dalla Norvegia e dalla Germania, si evidenzia come la sua accuratezza sia conforme allo stato dell'arte al fronte di costi computazionali ridotti. Lo sviluppo di algoritmi *machine learning*, supervisionati e non, di previsione dell'occupazione degli edifici con un'accuratezza superiore al 97%, aiuta a identificare le finestre più adatte per le opportunità di risparmio energetico e a fornire informazioni sui modelli di presenza e assenza. Inoltre, un algoritmo di automazione del riscaldamento e dell'aria condizionata basato su regole cerca di sbloccare l'enorme potenziale di risparmio energetico degli edifici senza compromettere il comfort termico degli occupanti. Le simulazioni su dati

dal Portogallo dimostrano la possibilità di avere un risparmio energetico superiore al 15%.

Spostando il focus dagli edifici alle comunità intelligenti, ci concentriamo sull'importante ruolo che la mobilità verde intelligente ha nel sostenere l'ulteriore digitalizzazione del settore dell'energia elettrica. Per superare gli inconvenienti dovuti alla scarsità delle infrastrutture di ricarica e facilitare l'adozione dei Veicoli Elettrici (EV), presentiamo un metodo EV specifico di *Reinforcement Learning* (RL) che garantisce l'itinerario energeticamente fattibile nel contesto della teoria dei grafi. Di conseguenza, proponiamo diversi algoritmi *deep*-RL per controllare la ricarica dei EV con l'obiettivo di aumentare l'autoconsumo delle energie rinnovabili e la soddisfazione dei conducenti. Le valutazioni effettuate rispetto al controllo basato su regole e modelli predittivi dimostrano le superiori prestazioni computazionali del RL e la migliore adattabilità ai futuri sistemi di mobilità. Infine, introduciamo un innovativo framework decentralizzato supportato dalla *blockchain* che consente una contabilità sicura e affidabile degli scambi di energia all'interno della comunità intelligente. La sua implementazione in Svizzera mostra il potenziale della *blockchain* nel ridurre i costi di ricarica di EV, trasformare il modello di business del mercato e aumentare la presenza di EV. Questa tesi si conclude con delle raccomandazioni rivolte ai principali stakeholder dell'energia al fine di far progredire ulteriormente la sua digitalizzazione.

Parole chiave: intelligenza artificiale, *blockchain*, digitalizzazione, veicolo elettrico, efficienza energetica, sistema di gestione dell'energia, tecnologia dell'informazione e della comunicazione, *machine learning*, fotovoltaico, *reinforcement learning*, edificio intelligente, comunità intelligente

Acknowledgements

During these years of my PhD thesis, I was fortunate to share my journey with so many knowledgeable, interesting, and talented people who motivated me daily to become the best version of myself and greatly contributed to my professional development. Therefore, I would like to express my deepest gratitude to those without whom this experience would not have been possible. First of all, I would like to thank my thesis supervisors, Christophe Ballif and Nicolas Wyrsh, for giving me the opportunity to conduct research in the PV-Lab and genuinely supporting my work. Especially, I would like to sincerely thank Nicolas for giving me the freedom to experiment, develop, and implement my own research ideas, continuously expanding my knowledge and sharpening my skills. Moreover, I would like to thank him for always being available for discussions, keeping up a great team spirit, even in the difficult times of teleworking, and sharing numerous work-related trips to different countries worldwide.

During the four years of my research, I was involved in FEEdBACK, the EU Horizon 2020 project that gave me a chance to collaborate and learn from so many brilliant people. I would like to express my gratitude to Filipe Soares and André Madureira for managing the project and facilitating efficient interactions within our team. I want to sincerely thank António Barbosa for his tremendous help in data management and releasing my algorithms into production. I want to thank Fernando Ribeiro and João Viana for scientific collaborations, Matthias Ansbach and Thekla Merfort for their efforts in data collection, and Andreu Pagès and Oriol Pla for their database support. Moreover, I would like to thank António Coelho, Arjan Van Timmeren, Chris Parker, Elisa Linares, Fernando Cassola, Ivan Capdevila, Jan Belger, Jesper Thestrup, John Dale, José Andrade, Joana Coelho, Klaus Schafmeister, Laura Martinez, Miguel Cruz, Nilufar Neyestani, Niall Castelli, Nelson Morais, Patricia Vale, Paul Ormerod, Rick Fransman, Sonia Roura, Trine Sørensen, and Tiago Santos for making FEEdBACK an enjoyable virtual workspace and creating great memories together. It was a pleasure to work with all of you.

Although the digitalization project I participated in was substantially smaller in size, it allowed me to greatly expand my knowledge in blockchain and acquire hands-on skills. Therefore, I would like to extend my gratitude to David Wannier, Jérémie Vianin, Jean-Marie Alder, and Vlado Mitrovic. Thank you for our profound discussions of the subject and useful workshops.

As a PhD candidate, I had a chance to supervise several semester/master projects and meet students from all over the globe. Particularly, I would like to thank Yann Martinson, Mahrok Ghoddousi, Loïc Martin, Sélim Kamal, and Xiaolong Lu for their hard work and our fruitful exchanges of research ideas.

Acknowledgements

I would like to sincerely thank my brilliant teammates Lionel, Jordan, and Marine for their continuous support, motivation, and inspiration. It was a pleasure to share my PhD journey with you. Thank you for making the systems group an enjoyable working environment.

Special gratitude goes to all amazing PV-Lab members, side by side with whom I had the opportunity to work all these years. In particular, I would like to thank Jean, Esteban, Luca G., Peter, Julie, and Franz for our great and occasionally challenging times spent teaching the PV fundamentals course. Moreover, I want to thank Philippe, Quentin J., and Xin Yu for being my office-mates, Karine for her highly appreciated administrative support, Hassan for his timely IT help, and Federica for being a like-minded person and a great supporter. I would also like to extend my gratitude to Aïcha, Alejandro, Alessandro, Ana, Andrea, Andrew, Daniel, Ezgi, Fabiana, Florent, Jonathan, Luca A., Mario, Olatz, Samira, Sofia, Quentin G., and many others for creating wonderful memories that I will cherish my whole life.

Continuing to express my gratitude, I would like to sincerely thank Rachid Cherkaoui for following my journey and being a great mentor and friend. Moreover, I want to thank Olivier van Cutsem whose thesis served as a source of inspiration during my writing of this manuscript.

Finally, I would like to sincerely thank my parents, my brother, and the rest of my family for their endless faith in me and their immense support during all these years. Despite the distance, your genuine celebration of my achievements and your example of hard work gave me the strength to keep going and set ambitious goals for myself. To conclude, I would like to wholeheartedly thank Alessandro, my incredible life partner and my biggest supporter, for living together through this memorable experience. Thank you for sharing my ups and downs, passionately discussing my scientific ideas, joining me on my travels worldwide, and cheering me up 24/7 during the last year and a half of my PhD spent in lockdown. Your love is my endless source of motivation, and without you, none of this would be even possible.

Marina Dorokhova

Prilly, 11 May 2021

Contents

Abstract	iii
Acknowledgements	xi
List of figures	xvi
List of tables	xix
Acronyms	xxv
1 Introduction	1
1.1 Background and motivation	1
1.1.1 Energy industry	2
1.1.2 Digitalization drivers	6
1.1.3 Advances in ICT	10
1.2 Thesis contribution	12
1.3 Thesis organization	15
2 Problem Formulation	17
2.1 Research objectives	17
2.2 Research methodology	18
2.2.1 Selected ICT technologies	19
2.2.2 Energy use cases	19
2.2.3 Energy goals	19
I Energy Management in Smart Buildings	21
3 The Smart Building ICT Framework	23
3.1 Background and motivation	24
3.2 The FEEdBAck project	24
3.3 The ICT-based smart building platform	25
4 Load Disaggregation	29
4.1 Background and motivation	30
4.2 State of the art	31
	xiii

Contents

4.3	Methodology	32
4.3.1	Device Usage Estimation algorithm	32
4.3.2	Benchmarking	35
4.4	Case studies	37
4.4.1	Norway	37
4.4.2	Germany	41
4.5	Conclusion	45
5	Occupancy Forecasting	47
5.1	Background and motivation	48
5.2	State of the art	49
5.3	Methodology	53
5.3.1	Supervised method	53
5.3.2	Unsupervised method	56
5.4	Case study	58
5.4.1	Simulation setup	58
5.4.2	Performance metrics	59
5.4.3	Baselines	59
5.5	Results	60
5.5.1	Supervised method	60
5.5.2	Unsupervised method	64
5.6	Conclusion	67
6	Building Automation	69
6.1	Background and motivation	70
6.2	State of the art	71
6.3	Methodology	72
6.4	Case study	74
6.4.1	Simulation setup	74
6.4.2	Performance metric	75
6.4.3	Baseline	75
6.5	Results	76
6.6	Conclusion	79
II	Energy Management in Smart Communities	81
7	The Smart Community Framework	83
7.1	Background and motivation	84
7.2	The Digitalization project	84
7.3	The ICT-based smart community framework	85

8	Routing of Electric Vehicles	89
8.1	Background and motivation	90
8.2	State of the art	91
8.3	Methodology	92
8.3.1	Environment	93
8.3.2	Markov decision process	94
8.3.3	Algorithm	96
8.4	Case study	99
8.5	Results	100
8.5.1	Training phase	100
8.5.2	Validation phase	102
8.5.3	Discussion	102
8.6	Conclusion	103
9	Energy Management of Electric Vehicles	105
9.1	Background and motivation	106
9.1.1	Control methods	106
9.2	State of the art	108
9.3	Methodology	110
9.3.1	Environment	110
9.3.2	Markov decision process	112
9.3.3	Algorithms	115
9.4	Case study	122
9.4.1	Datasets	122
9.4.2	Performance evaluation	123
9.4.3	Benchmark algorithms	124
9.4.4	Implementation	126
9.5	Results	126
9.5.1	Training phase	126
9.5.2	Testing phase	129
9.6	Conclusion	137
10	Blockchain Energy Exchanges	141
10.1	Background and motivation	142
10.1.1	Blockchain fundamentals	142
10.1.2	Blockchain glossary	145
10.2	State of the art	147
10.3	Methodology	148
10.3.1	Choice of blockchain	148
10.3.2	Smart contract	153
10.4	Implementation	155
10.4.1	Process flow	155
10.4.2	Web interface	156

Contents

10.4.3 Mobile application	158
10.5 Case study	160
10.6 Results	164
10.7 Conclusion	170
Conclusion	172
A Appendix	179
A.1 Characteristics of households	180
A.2 Characteristics of occupancy detection sensors	182
A.3 Feature engineering in occupancy forecasting	187
A.4 Taxonomy of reinforcement learning	188
Bibliography	189

List of Figures

1.1	Conventional energy value chain	2
1.2	Future energy value chain	3
1.3	Global LCOE from newly commissioned utility-scale renewable power generation technologies, 2010-2019. Source: [IRENA 2020]	5
1.4	Annual size of the global datasphere. Source: [Reinsel 2018]	6
1.5	Supercomputer power measured as the number of floating-point operations carried out per second. Source: [Roser 2020]	7
1.6	Expected growth of connected devices worldwide. Source: [NCTA 2020]	8
1.7	Key milestones in the history of AI	9
1.8	Technological drivers of digitalization	10
1.9	Gartner Hype Cycle for Digital Grid Transformation Technologies 2020. Source: [Sumic 2020]	12
2.1	Proposed energy systems research framework	18
3.1	The architecture of the FEEdBACK ICT-based platform	26
4.1	The Device Usage Estimation (DUE) algorithm explained	32
4.2	Example of a load curve disaggregated using DUE algorithm	33
4.3	<i>Energy Share Error</i> for all algorithms across all datasets	36
4.4	<i>Estimation Accuracy</i> versus execution time for all algorithms across three datasets	37
4.5	Breakdown of residential end-use electricity demand in Norway 2006 [Dalen 2013]	39
4.6	Load disaggregation example of one day in Norwegian apartment	40
4.7	Load disaggregation example of one day in German household ID1007	42
4.8	Breakdown of electrical energy consumption in Germany's private households 2006 [Schweiss 2006]	43
5.1	Taxonomy of occupancy forecasting based on input and output flows	52
5.2	Supervised pipeline for occupancy forecasting from electrical consumption data	53
5.3	Unsupervised pipeline for occupancy forecasting from ambient environment data	56
5.4	Example of unsupervised K-means clustering technique applied to indoor temperature measurements	57
5.5	Description of the majority voting procedure	57
5.6	The dataset distribution according to 'presence' and 'absence' classes	58

List of Figures

5.7	A 3D representation of labeled electrical consumption data obtained using PCA	60
5.8	Confusion matrices on test data for linear SVM, AdaBoost and feedforward ANN with dropout models	63
5.9	Grid search optimization of the linear SVM's prediction accuracy as a function of regularization parameter C	64
5.10	LSTM's epoch size optimization reported along with the model's execution time. All experiments were conducted with a fixed batch size of 32, which is a default value for ANNs in Keras Python library.	66
5.11	LSTM's batch size optimization reported along with the model's execution time. All experiments are conducted with fixed number of epochs equal to 10.	66
6.1	Occupancy-based HVAC scheduling model block-scheme	74
6.2	Simulation 1, October 2018 - March 2019: comparison between conventional and improved ON/OFF schedules of the HVAC system (segment's width = ON or OFF duration).	76
6.3	Simulation 2, July 2019 - September 2020: comparison between conventional and improved ON/OFF schedules of the HVAC system (segment's width = ON or OFF duration).	78
7.1	The architecture of the ICT-based smart community framework	86
8.1	The concept of reinforcement learning	92
8.2	Matrix representation of a combined state-action space	94
8.3	One-step ahead Q-learning algorithm	97
8.4	Route planning process	98
8.5	Graph representation of the road network of the Val d'Hérens alpine region in Switzerland	99
8.6	Example of a learning curve obtained during the training process. The bold line and the shaded region show the mean and the standard deviation of 5 runs. . .	100
8.7	The influence of binning parameter on the training time and state space's size. All experiments are conducted with a fixed number of episodes equal to 10000. .	101
9.1	Energy system representation	110
9.2	Representation of parametrized action space	113
9.3	Training performance of deep RL algorithms	127
9.4	Training duration of RL algorithms	128
9.5	Examples of benchmark and RLC algorithms on selected episodes of test data .	131
9.6	RLC and benchmark algorithms' comparison across episodes	132
9.7	Comparison of grid power usage across algorithms	132
9.8	Comparison of SOC evolution across episodes with equal SOC at arrival	133
9.9	Comparison of SOC evolution across episodes with binned charging duration .	135
10.1	Blockchain concept	143
10.2	Message encryption and digital signature concepts in blockchain	143

10.3 EV charging process flow	156
10.4 Energy flow tab on the web interface	157
10.5 Charging station management tab on the web interface	157
10.6 Clients tab on the web interface	158
10.7 Login page	159
10.8 Account page	159
10.9 Map of charging stations	159
10.10 Charging station details	159
10.11 Real-world test infrastructure	160
10.12 The "Add new vehicle" dialog window	161
10.13 A confirmation message at the registration of a new EV	161
10.14 The EV inventory of Hotel du Pigne	162
10.15 The record of transaction on Etherscan	163
10.16 The web interface of the Hotel Aiguille de la Tza	163
10.17 The web interfance of the Hotel du Pigne	164
10.18 The recommended gas price as of February 4th, 2021	165
10.19 The gas calculator	166

List of Tables

4.1	List of possible activities and related appliances	33
4.2	Appliances and their nominal power grouped per category	34
4.3	Testing and training periods for the three datasets	35
4.4	Appliances present in the ground truth of Norwegian demonstration households	38
4.5	Load disaggregation results of Norwegian households: ground truth and predicted <i>Energy Share</i> [%] per category followed by the <i>Energy Share Error</i> [%] per category	38
4.6	<i>Energy Share Error</i> [%] per category deduced from comparing load disaggregation results to Norwegian statistics of electricity end-use with and without heating	40
4.7	Load disaggregation results for German households: predicted <i>Energy Share</i> [%]	42
4.8	German demonstration households' median <i>Energy Share Error</i> across categories	43
4.9	Categories' <i>Energy Share</i> [%] comparison between GELAP dataset [Wilhelm 2021] and German demonstration dataset. The data is grouped according to the number of household's inhabitants and weekday/weekend feature.	44
5.1	Evaluation of feature selection methods for supervised occupancy forecasting models. On cross-validation, the mean accuracy and (standard deviation) across folds are reported, while on test, only the accuracy.	61
5.2	The impact of feature engineering on the results of supervised occupancy forecasting. On cross-validation, the mean accuracy and (standard deviation) across folds are reported, while on test, only the accuracy.	62
5.3	Accuracy and (standard deviation) of unsupervised occupancy forecasting models	65
6.1	Distribution of potential energy savings among seasons	78
6.2	Distribution of potential energy savings among HVAC functionalities	79
6.3	Evaluation of potential energy savings in a university building in Portugal	79
9.1	Characteristics of chosen RL algorithms	116
9.2	Hyperparameters of the DDQN algorithm	117
9.3	Hyperparameters of the DDPG algorithm	119
9.4	Hyperparameters of the P-DQN algorithm	121
9.5	Share of episodes with various SOC at departure satisfaction levels	134
9.6	Total PV self-consumption and total energy purchased from the grid	136

List of Tables

9.7	The algorithms' execution time on test dataset	136
10.1	Qualitative comparison of various blockchain frameworks	149
10.2	Assessment comparison of various blockchain frameworks	151
10.3	The charging transaction object	153
10.4	The data collected during the experiment	166
10.5	Case study data	168
10.6	Charging-event related costs for EV driver in [CHF]	168
A.1	Characteristics and appliances' inventory of Norwegian demonstration households	180
A.2	Characteristics and appliances' inventory of German demonstration households	181
A.3	Characterization of sensor technologies used for occupancy detection	182
A.4	Advantages and disadvantages of sensor technologies used for occupancy detection	184
A.5	Sets of features used to transform raw data in supervised occupancy forecasting	187

Acronyms

AI	Artificial Intelligence
ANN	Artificial Neural Network
API	Application Programming Interface
AR	Augmented Reality
ARIMA	Auto Regressive Integrated Moving Average
BLE	Bluetooth Low Power iBeacon
BMS	Building Management System
CO	Combinatorial Optimization
CSO	Charging Station Operator
DAO	Decentralized Autonomous Organization
DDPG	Deep Deterministic Policy Gradient
DDQN	Double Deep Q-Networks
DDSC	Discriminative Disaggregation via Sparse Coding
DQN	Deep Q-Network
DSM	Demand-Side Management
DUE	Device Usage Estimation
EMSP	Electric-Mobility Service Provider
ESS	Energy Storage System
EU	European Union
EV	Electric Vehicle
EWC	Energy Web Chain
FEEdBACK	Fostering Energy Efficiency and Behavioral Change through ICT
FHMM	Factorial Hidden Markov Model
FN	False Negative
FP	False Positive

Acronyms

GDPR	General Data Protection Regulation
GHG	Greenhouse Gas
GPS	Global Positioning System
GSP	Graph Signal Processing
GUI	Graphical User Interface
HER	Hindsight Experience Replay
HMM	Hidden Markov Model
HVAC	Heating, Ventilation, and Air Conditioning
ICT	Information and Communication Technology
IoT	Internet of Things
IT	Information Technology
KNN	K Nearest Neighbours
LCOE	Levelized Cost of Electricity
LSTM	Long Short-Term Memory
MARL	Multi-Agent Reinforcement Learning
MDP	Markov Decision Process
MILP	Mixed-Integer Linear Programming
ML	Machine Learning
MPC	Model Predictive Control
NILM	Non-Intrusive Load Monitoring
P-DQN	Parametrized Deep Q-Networks
PAYG	Pay As You Go
PBFT	Practical Byzantine Fault Tolerance
PC	Personal Computer
PCA	Principal Component Analysis
PIR	Passive Infrared Sensor
PoA	Proof-of-Authority
PoB	Proof-of-Burn
PoS	Proof-of-Stake
PoW	Proof-of-Work
PV	Photovoltaic
RBC	Rule-Based Control

ReLU	Rectified Linear Unit
RES	Renewable Energy Sources
RFECV	Recursive Feature Elimination with Cross-Validation
RFID	Radio-Frequency Identification Tag
RH	Relative Humidity
RL	Reinforcement Learning
RLC	Reinforcement Learning Control
SCCER	Swiss Competence Center for Energy Research
SOC	State of Charge
SVM	Support Vector Machine
TD	Temporal Difference
TN	True Negative
TP	True Positive
TPS	Transaction Per Second
TSDB	Time Series Database
UWB	Ultra-Wideband Sensor
V2G	Vehicle-to-Grid
VR	Virtual Reality
WLTP	Worldwide Harmonised Light Vehicles Test Procedure

1 Introduction

1.1 Background and motivation

Industry 4.0 is the fourth major transformation happening to humankind since the First Industrial Revolution of the 18th century. Building on the foundations of three preceding revolutions, enabled by mechanization, electrification, and computerization, Industry 4.0 is destined to blur the boundaries between physical, digital, and biological worlds [Schwab 2016]. Characterized by emerging technological breakthroughs in a number of different research fields, such as computer science, robotics, and telecommunications, it is fundamentally changing the way we live, work, and interact with each other.

It is widely argued that the Fourth Industrial Revolution will be unlike anything the world has seen before. Being unprecedented in its speed, size, complexity, and scope, it evolves rather exponentially than linearly, shifting paradigms in science, economy, business, and politics. Contrary to previous worldwide industrial transformations brought by steam engines, electric light bulbs, and personal computers, the modern era's technologies are more sophisticated and pervasive in nature. Moreover, the products that stem from them emerge and mature at much higher, extraordinary speeds: the adoption of landline telephones took almost 50 years, while the internet and Twitter became our everyday commodities in only 7 years and 9 months, respectively [Desjardins 2018].

The deep fusion of new developments that leverage the pervasive power of Information and Communication Technology (ICT) affects all sorts of systems across key societal functions. Despite the discrepancy in the timeline, with 13% of the world's population still lacking access to electricity nowadays [Ritchie 2019], one of the merits of the Second Industrial Revolution, the Fourth Industrial Revolution will inevitably penetrate every country, industry, and institution.

Being one of the essential sectors of the world economy, the energy industry is widely and profoundly affected by fast-moving changes, which implies both - finding efficient solutions for existing challenges and creating new ones to overcome. So why does energy need the transformations brought by the Fourth Industrial Revolution?

1.1.1 Energy industry

The energy industry and the electric power sector, in particular, are going through challenging times of extensive transformation and disruptive changes, finding themselves at the crossroads of the three "D"s: decarbonization, decentralization, and digitalization [Di Silvestre 2018a]. These three phenomena accelerate the countdown to a new energy world faster than expected, calling the energy economy's key stakeholders for immediate action.

The urge for decarbonization goes beyond the volatile prices of fossil fuels and energy dependencies between countries and is primarily explained by the dramatic deterioration of the world's climate that endangers the future of humanity. The widespread adoption of Renewable Energy Sources (RES) and improvements in energy efficiency related to power generation, transport, and usage contribute to shifting away from fossil fuels. However, drastic changes to international and national policies are required to address climate change through imposing ambitious targets on CO₂ emissions reduction. The European Union (EU) has recently proposed to aim at 55% emissions reduction by 2030, thus aspiring to limit the global temperature rise to 1.5°C and to become climate-neutral by 2050 [EU 2020]. While decarbonization is merely an umbrella term that encompasses various instruments to combat climate change, decentralization and digitalization will play a key role in achieving a carbon-free future.

The decentralization is promoted by the change in the power production and consumption paradigm, where a more distributed generation is used to satisfy the energy needs. The concept is tightly linked to decarbonization as the majority of the distributed generators are renewable in nature. Being primarily designed for transporting energy from source to load using electricity, the energy system has been successfully coping with delivering energy services to the customers over the past years. Its conventional structure, shown in Figure 1.1, consists of several successive elements. The energy value chain starts with power generation, whose sources range from fossil fuel combustion and nuclear fission towards all kinds of RES, including solar, wind, hydro, and geothermal. Once electricity is generated, the transmission grid carries it at a high-voltage over long distances until the end-user proximity is reached. Passing through the step-down transformer, thereby lowering its voltage level, electricity is consecutively delivered by the distribution network with the help of utility providers to residential and commercial customers, ready for everyday use.



Figure 1.1 – Conventional energy value chain

The described above structure represents the electrical grid designed at the dawn of the 19th century, when the initiation of electrification was characterized exclusively by top-down power flows, fossil fuel generation, and centralized governance. However, nowadays, the reality of increased RES penetration and availability of advanced technologies are reshaping the energy value chain, as seen in Figure 1.2. The emergence of new energy actors, prosumers, whose functionality goes beyond energy consumption towards energy production and trading, enables a transition from the legacy centralized power network to the next-generation decentralized electrical grid.




Figure 1.2 – Future energy value chain

The energy industry's digitalization is driven by the rapid development of Artificial Intelligence (AI) and advances in ICT characterized by an increased number and volume of data exchanges. The electric grid's ever-growing complexity requires new management approaches and new business models to empower efficient interactions between consumers, prosumers, grid operators, energy providers, and legal authorities. Being the core concept of the Fourth Industrial Revolution, digitalization links together digitization and digital transformation processes. The formal definition of digitalization is given by [Santamaria 2019]:

Digitalization

The use of digital technologies to change a business process and enhance efficiency and revenue; it is the process of moving to a digital business.

Therefore, the above-mentioned terms form a digital hierarchy of Industry 4.0. Digitization converts the information from analog to digital form, building foundations for digitalization to further change existing business processes using digital technologies. The digital transformation goes beyond digitalization and results in profound changes in business models, creation of new markets, revenue streams, and value-producing opportunities. While the energy industry's digitization is ongoing successfully, embracing the next level, digitalization, requires substantial transformations across the energy value chain. However, such extensive changes come with considerable benefits: digitalization is expected to provide a 25% reduction in operational expenses and 20% to 40% performance gains in safety, reliability, and customer satisfaction [McKinsey 2018]. Moreover, the joint influence of decarbonization, decentralization, and digitalization fosters the energy transition and calls for disruptive approaches to resolve the energy industry's prominent challenges as identified in [Attwood 2017]:

- **Intermittent nature of renewable energy** complicates energy supply forecasting and increases the demand for balancing services due to a mismatch between sources availability and established consumption patterns. Failure to consistently deliver energy can disturb the grid's frequency, voltage level, and pose significant threats to the electric grid's stable, reliable, and secure operation. The AI-enabled digital instruments can raise the energy yield of renewables and accurately predict demand and supply, thus increasing the dispatching efficiency, improving the grid's reliability, and reducing the need for balancing reserves. For instance, real-time sensors' information enhanced by Machine Learning (ML) boosts the energy production of wind farms by 20% [Bughin 2017], while self-learning weather models coupled with intelligence applied to large historical datasets and instantaneous measurements from Photovoltaic (PV) installations improve solar forecasting by 30% [Mortier 2020].
- **Aging electrical grids**, built in developed countries more than 50 years ago, do not correspond anymore to the customer needs in terms of transmission and distribution lines capacity. A limited number of options to solve this problem, namely construction of new lines or energy curtailments, challenges possible grid refurbishment strategies. The AI-enabled planning of grid operations can prolong the critical assets' life using preventive maintenance, automate fault prediction, and decrease downtime. Moreover, it can reduce transmission and distribution energy losses, eliminate the need for substantial capital investments into new power infrastructure, and even increase network operators' earnings by 20% to 30% [Bughin 2017].
- **Growing number of independent power producers** increases overall system complexity and demands new managing strategies to effectuate reliable energy services. The exclusive rights of utilities for owning generation assets are questioned, making them shifting their primary focus towards managing energy supply. Despite diversifying power sources at the edge of the grid improves the grid's resilience towards outages and blackouts, the safe and sound operation of a large number of independent generators requires process optimization and enhanced connectivity. The future energy providers can facilitate efficient integration of distributed generators into the broader smart grid by leveraging tamper-proof blockchain-enabled smart meters and intelligent AI-supported decision-making for real-time optimal power dispatch and peak demand anticipation.
- **Increased self-consumption** represents the transformation of consumers into prosumers, which can not only satisfy their own needs but also provide an excess of generated power to the market and ancillary services to the grid. This requires a new economic framework to settle the competition, where each of the participants will successfully secure their revenue streams. The AI-powered digital platforms can intelligently match power sellers and buyers and efficiently utilize a growing number of smart meters to enable ML-based dynamic pricing. Moreover, the leverage of blockchain technology can empower energy communities with new-level relationships formed on the principles of security, privacy, immutability, and efficient conflict-resolution.

1.1. Background and motivation

The challenges the power sector is facing did not arise overnight, and the influence of the three "D"s on the current operations and future development of electrical grids proved itself to be irreversible. Fostered by technological improvements, economies of scale, increasingly competitive supply chains, and growing developer experience, the Levelized Costs of Electricity (LCOEs) from key clean energy technologies continue to decline as seen in Figure 1.3.




Figure 1.3 – Global LCOE from newly commissioned utility-scale renewable power generation technologies, 2010-2019. Source: [IRENA 2020]

The PV technology, whose LCOE has seen a dramatic decrease of 82% between 2010 and 2019, takes the lead among other RES thanks to continued improvements in module efficiency and modernization of manufacturing processes, which lower material costs and raise productivity through extensive automation [IRENA 2020]. Accompanied by drastically decreasing prices of Energy Storage Systems (ESSs), with lithium-ion battery packs projected to reach \$100/kWh by 2023 (-85% from 2013) [Henze 2020], PV generation at or below \$20/MWh followed by wind at \$30/MWh in a good site is the cheapest source of electricity in history [IEA 2020b] and claims major responsibility for changing the energy production and consumption paradigm in the energy industry. With LCOEs of key RES confidently declining and matters of climate change becoming more pressing, the accelerated penetration of renewables becomes inevitable. Thus, the future grid has no other choice but to evolve into a robust and adaptive infrastructure capable of playing by the rules of the Fourth Industrial Revolution and accommodating new and emerging technologies. The integration of power systems with ICT will signify a new era in the energy industry and unlock tremendous opportunities for the power sector. Ultimately, the ubiquitous digitalization may lead to entirely independent electrical grids where the operations are optimized automatically and decisions are made without any human intervention. However, why is now the best time for immediate action and what are the key drivers and technologies that empower digitalization in energy?

1.1.2 Digitalization drivers

Although the birth of AI dates back to the middle of the 20th century, the digital revolution has taken off only recently due to the emergence of powerful technological trends. According to [Attwood 2017, IEA 2017, Bughin 2018] one can distinguish four main factors that fuel digitalization: explosion of data, improvements in computing power and capacity, widespread coverage of communication networks, and progress in the field of ML. While each of these drivers represents a trend that grows exponentially, their simultaneous influence creates a real momentum for digitalization to propagate through industries and change them once and forever.

Explosion of data

The importance of data, commonly being framed as the new oil of the digital economy, has been growing steadily over the past decade, given the success of data-driven decision-making and the contribution of data-derived insights to the knowledge base. Similar to oil, raw data holds little value and requires thorough processing to create valuable products. However, treating data in a timely and secure manner becomes increasingly difficult due to its exponential growth, as seen in Figure 1.4. According to [Reinsel 2018], the size of the global datasphere is projected to reach 175 ZB (10^{21}) by 2025: if we were able to store such amounts of data on DVDs, stacking them would be sufficient to circle the Earth 222 times.




Figure 1.4 – Annual size of the global datasphere. Source: [Reinsel 2018]

In the energy industry, the growing volume of data is driven by digitization, which results in extensive deployment of smart meters and sensors across various energy assets. However, the volume is not the only distinguishing characteristic of energy big data and is accompanied by velocity, variety, and value [Zhou 2016]. Velocity refers to the speed of collecting and processing data and deploying data-generated insights, which becomes critical when switching to real-time operations and decision-making. Variety is determined by the availability of structured and unstructured data, with the prior prevailing currently in the power sector. Finally, the value of energy big data is sparse, which requires mining for knowledge and insights across massive amounts of data using sophisticated algorithms and data processing techniques.

It is not an understatement to say that data is a cornerstone of digitalization as implementing AI-powered decision-making, exploiting self-learning systems, and leveraging advanced analytics for enhancing the efficiency of operations is not possible without data. However, possessing an abundance of digitized historical and real-time information requires significant efforts to turn it into a valuable asset. First, one has to ensure that specific pipelines are put in place to effectuate skillful data management, thus collecting, storing, processing, and analyzing data. Second, coming from various sources, the growing amounts of data demand dedicated efforts towards cybersecurity. Third, data ownership, handling, and governance guidelines, which are particularly relevant for energy utilities, have to be established. A successful combination of the above measures unleashes the power of digitalization and ensures that data-enabled operations respect privacy, security, and ethics.

Increasing computing power

The complexity of integrated circuits, governed by Moore's Law, has been doubling regularly since 1975, thus enabling faster and cheaper computers [Roser 2020]. The increasing trend in processing power, depicted in Figure 1.5, facilitates more efficient value creation through the possibility to handle larger volumes of data at shorter times. Advances in computing capabilities became the key enabler of technological breakthroughs, opening the window of opportunities for more powerful and sophisticated analytics, modern ML algorithms, and blockchain. Despite the confident and predictable growth in processing power, the future scenarios suggest that Moore's Law progression cannot be sustained indefinitely and will see its end by 2025 [Waldrop 2016]. The ever-decreasing in size microprocessors' features are approaching physical limits of feasibility, thus posing new challenges for the semiconductor industry to find a suitable successor for silicon-based technology and keep up with satisfying demands of the digitalization era.




Figure 1.5 – Supercomputer power measured as the number of floating-point operations carried out per second. Source: [Roser 2020]

Connectivity growth

Since the advent of the internet in 1991, the way how we communicate with each other has changed completely, allowing humanity to embrace instantaneous connections that see no geographical borders. Further development of mobile internet, that nowadays accounts for 51% of the website traffic worldwide [Clement 2020], has liberated society from being associated with one specific device and gave a possibility to remain constantly online. Despite the rapid expansion of internet technology, its geographical coverage is far from universal, with nearly 40% of the world's population still lacking access to internet connection [Grijpink 2020]. However, heterogeneity in distribution did not hinder the development of new types of interactions, such as human-to-machine and machine-to-machine. The latter gave rise to the phenomena of the Internet of Things (IoT), enabling communication and data exchange across devices in fully integrated networks.

The exponential growth of the number of connected devices worldwide, depicted in Figure 1.6, has far surpassed the size of the world's population and is projected to reach 150B by 2025 [Reinsel 2018]. By 2030, largely thanks to the development of the 5G internet, machine-to-machine interactions will represent more than 70% of all connections, generating 12% of the real-time data traffic worldwide [Grijpink 2020]. With faster speed and lower latency, the 5G creates opportunities to expand digitalization use cases and provide new services and applications. In the power sector, massive deployment of connected sensor networks encourages decentralization by allowing remote monitoring and control of energy assets. Moreover, connectivity growth facilitates real-time analytics to improve customers' understanding, helps to anticipate unexpected disturbances and faster locate faults in transmission and distribution networks, and enables new energy management solutions through instantaneous energy usage and pricing data.




Figure 1.6 – Expected growth of connected devices worldwide. Source: [NCTA 2020]

Advances in AI

The abundance of data, powerful computing, and connectivity growth have significantly influenced the world over the last decade. However, the switch to a digitalization mindset would not be possible without AI that aims to enhance and eventually replace the human's ability to analyze, manage, optimize, and make decisions intelligently. Since the middle of the 20th century, AI has come a long way, as seen in Figure 1.7, to become the main tool for solving the world's most prominent problems.

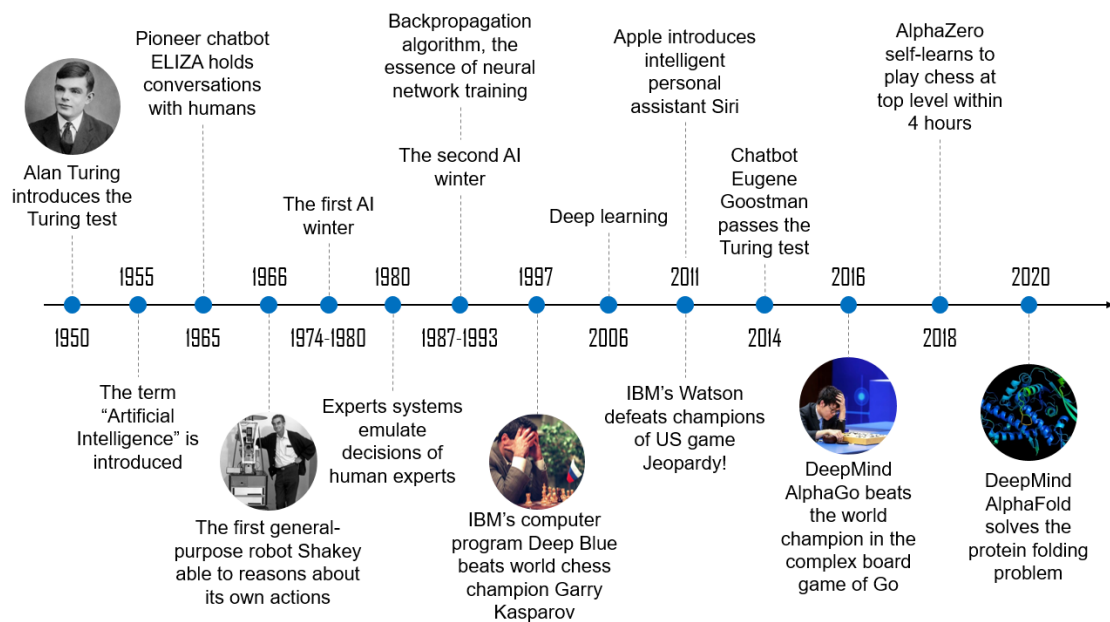


Figure 1.7 – Key milestones in the history of AI

After the second AI winter, ML, the subset of AI which makes computers capable of learning without explicit programming, began to flourish. In 2006 ML research has entered the deep learning era marked by the endpoint in the human-machine competition with DeepMind's AlphaGo defeating the 18-time world champion Lee Sedol in the complex board game of Go in 2016 [Cheng 2019]. Since then, AI has attained relevance in many different fields of our lives, with ML used to provide data-driven analytics, forecasts, control, and decision-making.

ML taxonomy can be divided into supervised learning, applied to regression and classification problems, unsupervised learning, used for clustering and dimensionality reduction, and Reinforcement Learning (RL). The latter in combination with multilayer Artificial Neural Networks (ANNs) represents a new frontier in AI research, where DeepMind's MuZero [Schrittwieser 2020] has recently made a significant step forward by mastering games of Go, chess, shogi, and Atari without prior knowledge of the rules. Tapping into the enormous potential of AI, there is no doubt how much it will affect the future of every industry and every human being. The question is: how to harness the power of AI to create more good?

1.1.3 Advances in ICT

With AI being the protagonist of digitalization, the influence of other key ICT technologies on the Fourth Industrial Revolution can be easily overlooked. Thus, Figure 1.8 summarizes the staple technological developments of the digitalization era, which beyond previously discussed AI, IoT, and big data, include cloud computing, edge computing, Augmented Reality (AR) and Virtual Reality (VR), blockchain, and cybersecurity [Attwood 2017, Faheem 2018].

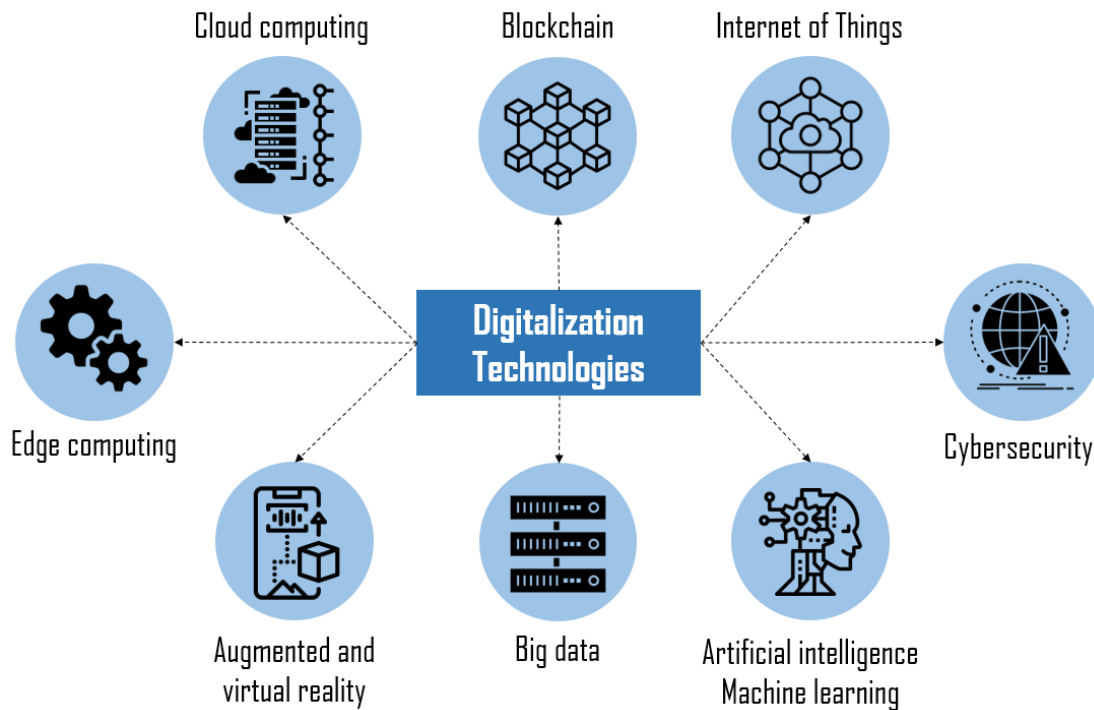


Figure 1.8 – Technological drivers of digitalization

- **Cloud computing** is the internet-enabled on-demand availability of computing services, which are distributed among a large number of data centers to achieve coherence and economies of scale. Thus, it eliminates the need to invest in the Information Technology (IT) infrastructure while allowing access, storing, and processing data using software applications remotely. For energy utilities embracing cloud computing would mean jump-starting digitalization through reductions in IT costs, security in storage of large amounts of data, access to best practices in advanced analytics and automation intelligence, and greater flexibility in asset management and optimization [Habibi 2013].
- **Edge computing** is a distributed computing concept that brings data processing closer to a location where the information is produced, thus allowing better bandwidth availability and improved response times. Particularly suitable for IoT technology, edge computing in the power sector would foster decentralization by streamlining the operation of distributed RES and unleashing the potential of advanced metering infrastructure.

- **AR and VR** refer to technological concepts of digitally enhancing the physical reality, with prior superimposing the digital information upon the real world while the latter providing a fully immersive experience. Although the number of implemented use cases is currently low, the adoption of AR and VR in the energy sector holds a great promise to increase operational safety, improve the speed and efficiency of maintenance activities, and enable better training of the workforce. With a lack of experienced fieldworkers in the energy industry, the latter becomes important for developing skills and knowledge while effectively responding to the challenge of limited training exacerbated by high cost and time investments, increased safety risks, and low accessibility of remote sites.
- **Blockchain** is a distributed ledger technology that allows information storage in the form of blocks of immutable and cryptographically-interlinked transactions. In the power sector, blockchain is seen as a revolutionizing technology that can encourage decentralization through various application use cases, such as peer-to-peer energy trading, Electric Vehicle (EV) charging, community energy storage, automated billing, real-time grid infrastructure management, and trading of renewable energy certificates. From the bigger perspective, the energy industry's adoption of blockchain is expected to improve energy security, reduce costs through process optimization, and promote sustainability by integrating higher share of RES in the energy mix [Andoni 2019].
- **Cybersecurity** refers to a combination of technologies, processes, and practices that aim to protect the IT infrastructure from malicious cyber attacks. The electric power sector is particularly vulnerable to cyber-threats, such as billing fraud, data theft, and ransomware, due to the large geographical spread, high organizational complexity, and tight interconnection with physical assets [Bailey 2020]. Being classified as critically important for society's functioning, electrical infrastructure is susceptible to cyber-attacks across the whole value chain, making cybersecurity paramount in the era of digitalization.

Characterized by different maturity levels, these technologies vary in their application scope, readiness for implementation, and potential to impact the power sector's development. Gartner Hype Cycle is the analysis tool that helps to understand how relevant the technology is for solving the industry's real problems and how its influence is expected to evolve with time. It consists of five areas that correspond to key phases of the technology's life cycle spanning from technological breakthrough and proof-of-concept to mainstream adoption and market applicability. Therefore, using the Gartner Hype Cycle, one can understand the interconnection between various technologies, evaluate incurred risks, and develop specific technology implementation strategies. Figure 1.9 represents the Gartner Hype Cycle of electrical power sector's emerging technologies that are paving the way for the industry-wide digital transformation. Although the majority of the aforementioned digitalization drivers are currently at the peak of inflated expectations or at the trough of disillusionment, the rapid progress of complementary technologies, such as EVs, smart meters, and ESSs, is expected to facilitate the adoption of ICT in the energy industry.

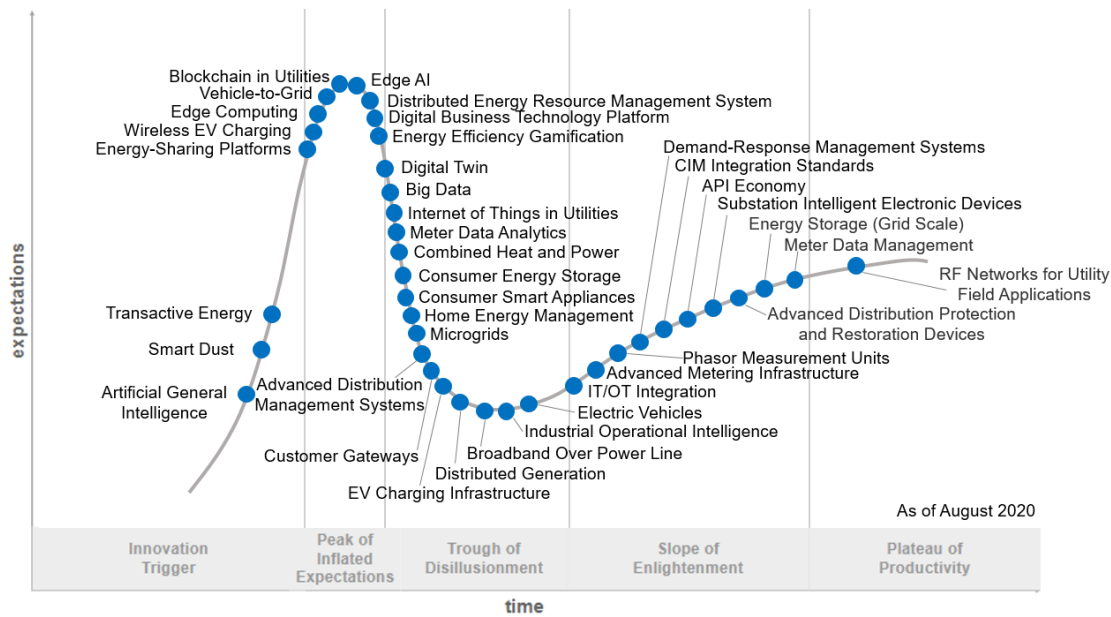


Figure 1.9 – Gartner Hype Cycle for Digital Grid Transformation Technologies 2020. Source: [Sumic 2020]

With the energy transition underway, it becomes evident that digitalization is the only way to keep the decarbonized and decentralized power sector up and running in a stable, secure, and affordable manner. However, when it comes to implementing profound changes to let the ICT shine, a large-scale digitalization is hindered by the industry's massive inertia. In a fast-developing world, where technologies emerge, progress, and become obsolete rapidly, most energy utilities are missing the real momentum by contemplating complex engineering plans to digitalize the industry instead of moving fast and decisively. With appropriate drivers being in place, now is the time to take bold actions and commit the power sector to take digital seriously.

1.2 Thesis contribution

This thesis aims to bridge the gap between strategic roadmaps focused on the energy industry's digitalization and their actual implementation in real-world scenarios through digital energy services. To address this challenge and help the electric power sector and its key stakeholders to navigate towards higher energy efficiency in the midst of the energy transition, a set of ICT-supported software applications, models, and tools that operate on top of existing physical infrastructure is proposed. Specifically, the research conducted in this thesis unfolds within the designed ICT-based smart building and smart community frameworks, the modular structure and scalability of which can serve as a backbone for future implementations of various digital energy management solutions and aid greatly the exploration of the digitalization phenomena in the energy industry.

While detailed contributions of the current thesis are featured at the beginning of each chapter, the main highlights of the conducted work can be summarized as follows:

- [1] The introduction of a novel unsupervised load disaggregation approach that helps to understand what drives the energy usage in residential households and raise the awareness of one's energy-related behavior without compromising privacy and security.
- [2] The development of ML-based supervised and unsupervised highly accurate occupancy forecasting algorithms that help to identify best-suited windows for energy-saving opportunities and deliver insights on one's presence and absence patterns.
- [2] The suggestion of an occupancy-centric rule-based Heating, Ventilation, and Air Conditioning (HVAC) automation algorithm that strives to unlock the buildings' massive potential for energy savings without compromising the occupants' thermal comfort.
- [6] The presentation of an EV-specific routing method based on RL techniques that aims to facilitate the large-scale deployment of electric mobility through overcoming the current problems related to sparse charging infrastructure and limited cruising range of EVs.
- [7] The development of several deep RL-based control algorithms for EV charging that ensure the smooth integration of electric mobility into the broader energy system through increasing the self-consumption of renewables and satisfaction of EV drivers.
- [8] The proposal of an innovative decentralized blockchain-supported framework that enables secure and reliable accounting of respective energy flows between the entities participating in the EV charging process in a cost-efficient and user-friendly manner.

Notably, the proposed tools and software applications resort to cutting-edge techniques and algorithms in the field of ICT, thus broadening the spectrum of possible methodologies to solve the energy industry's most prominent challenges.

The research conducted in this thesis led to the following publications of peer-reviewed articles and various presentations at scientific conferences and events, all of which are listed below:

Publications

- [1] Jordan Holweger, Marina Dorokhova, Lionel Bloch, Christophe Ballif and Nicolas Wyrsh. *Unsupervised algorithm for disaggregating low-sampling-rate electricity consumption of households*. Sustainable Energy, Grids and Networks, vol. 19, page 100244, 2019. DOI: 10.1016/j.segan.2019.100244
- [2] Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Rule-based scheduling of air conditioning using occupancy forecasting*. Energy and AI, vol. 2, page 100022, 2020. DOI: 10.1016/j.egyai.2020.100022

- [3] Albert Hoffrichter, Evangelos Zacharis, Angelina Katsifaraki, Ashley Morton, Gloria Calleja, Federica Fuligni, Marko Batic, Marina Dorokhova, Niall Castelli, Konstantinos Kanellos and Thanh Nguyen. *Behavioral Change towards EE by Utilizing ICT Tools*. MDPI Proceedings of The 8th Annual International Sustainable Places Conference (SP2020), vol. 65, no. 1: 10, 2020. DOI: 10.3390/proceedings2020065010
- [4] Filipe Soares, André Madureira, Andreu Pagès, António Barbosa, António Coelho, Fernando Cassola, Fernando Ribeiro, João Viana, José Ricardo Andrade, Marina Dorokhova, Nélson Morais, Nicolas Wyrsh, Trine Sørensen. *FEEdBACK: An ICT-based platform to increase energy efficiency through buildings' consumer engagement*. Energies, 14(6), 1524, 2021. DOI: 10.3390/en14061524
- [5] Marina Dorokhova, Fernando Ribeiro, António Barbosa, João Viana, Filipe Soares, Nicolas Wyrsh. *Real-world implementation of an ICT-based platform to promote energy efficiency*. Energies, 14(9), 2416, 2021. DOI: 10.3390/en14092416
- [6] Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Routing of electric vehicles with intermediary charging stations: A reinforcement learning approach*. Frontiers in Big Data, 4:586481, 2021. DOI: 10.3389/fdata.2021.586481
- [7] Marina Dorokhova, Yann Martinson, Christophe Ballif and Nicolas Wyrsh. *Deep reinforcement learning control of electric vehicle charging in the presence of photovoltaic generation*. Submitted, 2021
- [8] Marina Dorokhova, Jérémie Vianin, Jean-Marie Alder, Christophe Ballif, Nicolas Wyrsh and David Wannier. *A blockchain-supported framework for charging management of electric vehicles*. Submitted, 2021

Presentations

- 1. Marina Dorokhova, Jordan Holweger, Lionel Bloch, Christophe Ballif and Nicolas Wyrsh. *Unsupervised algorithm for disaggregating low-sampling rate electricity consumption of households*. Presented at EU NILM 2018: the 5th European Non-Intrusive Load Monitoring Workshop, Duisburg, Germany, 2018
- 2. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Supervised model for building occupancy prediction from electrical consumption data*. Presented at Applied Machine Learning Days, Lausanne, Switzerland, 2019
- 3. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Achieving savings through energy monitoring, forecasting and optimization: The European FEEdBACK project*. Presented at World Smart Energy Summit, Moscow, Russia, 2019
- 4. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Occupancy forecasting and load disaggregation: an energy efficiency perspective*. Presented at SEST 2019: the 2nd

International Conference on Smart Energy Systems and Technologies, Porto, Portugal, 2019

5. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Rule-based optimal HVAC scheduling using occupancy forecasting: A machine learning approach*. Presented at GYSS 2020: the Global Young Scientists Summit, Singapore, Singapore, 2020
6. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Energy-efficient routing of electric vehicles with intermediary charging stations: A reinforcement learning approach*. Presented at Applied Machine Learning Days, Lausanne, Switzerland, 2020
7. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Technological and legal challenges in the development of ICT-based FEEDBACK platform*. Presented at SP 2020: the 8th Annual International Sustainable Places Conference, online, 2020
8. Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Digitalization of the electric power sector: Machine learning-based applications for energy management*. Presented at WiDS 2021: the Women in Data Science Worldwide Conference, online, 2021

1.3 Thesis organization

This thesis manuscript is divided into two main parts. The first part deals with the energy management within the smart building, exposing the application of various AI-powered models to specific energy use cases. The second part revolves around the integration of smart buildings and EVs into the broader entity of the smart community, supported by AI and blockchain technologies. The chapters are organized as follows:

- **Chapter 2** defines the research methodology, particularly focusing on the selected ICT background, specific energy use cases, and desired energy objectives.

PART I - Energy Management in Smart Buildings

- **Chapter 3** presents an ICT-based platform allowing to explore digital energy management solutions at the smart building level. Through various interconnected software applications that operate on top of existing infrastructure, the platform aids interested stakeholders to optimize energy consumption without compromising comfort levels. Some particular instances of implemented digital energy services are presented in the following Chapters 4, 5, and 6.
- **Chapter 4** introduces a novel low-sampling-rate unsupervised load disaggregation algorithm based on activity-chain reconstruction through the Markov process. The algorithm's benchmarking to state-of-the-art methods in the field proves that it achieves similar disaggregation accuracy while exhibiting lower computational cost. The algorithm's performance is showcased on real-world datasets from Norway and Germany.

- **Chapter 5** proposes two highly-accurate context-based approaches to forecast binary occupancy using advanced ML techniques. The supervised method predicts occupancy from electrical energy consumption data, while the unsupervised one uses ambient environment measurements. The algorithms' performance is evaluated on a real-world dataset collected in Portugal.
- **Chapter 6** suggests a novel occupancy-centric rule-based automation algorithm for Heating, Ventilation, and Air Conditioning (HVAC). The methodology generates intelligent day-ahead ON/OFF HVAC schedules aiming to deliver energy savings. Evaluating the algorithm's performance on two real-world datasets collected in Portugal highlights the higher energy-saving potential of cooling as opposed to heating.

PART II - Energy Management in Smart Communities

- **Chapter 7** presents an ICT-based framework that brings together physical and digital assets to streamline the operations of the smart community. Particularly, it focuses on the important role of intelligent green mobility in supporting further digitalization of the electric power sector. Some EV-specific digital energy services that comprise the suggested framework are described in details in the following Chapters 8, 9, and 10.
- **Chapter 8** introduces an EV-specific routing method based on RL that aims to generate energy feasible paths from starting point to destination while considering both recharging possibilities at intermediary charging stations and the ability of EVs to recuperate energy. Formulated in the graph-theoretical context, the suggested approach is validated on the segment of a real-world road network located in Switzerland.
- **Chapter 9** proposes the methodology to enable deep RL control of EV charging with the objective to simultaneously maximize PV self-consumption and EV's state of charge at departure. Differing by the type of action space used, the 3 suggested algorithms are benchmarked against popular control methods in the field. The results demonstrate the RL's superior computational performance and better fitness for future mobility systems.
- **Chapter 10** describes a blockchain-based EV charging digital energy service that enables secure and reliable accounting of energy exchanges within the smart community through a specifically designed smart contract. Tightly interlinked with physical infrastructure and implemented in a real-world demonstration site in Switzerland, the suggested service demonstrates blockchain's cost-reduction potential and facilitates the large-scale deployment of electric mobility.

Finally, a **Conclusion** chapter summarizes the findings of this thesis, provides recommendations to the energy industry's key stakeholders, and further discusses the future outlook of the ICT's role in the digitalization of the energy industry.

2 Problem Formulation

2.1 Research objectives

With the energy value chain evolving, subject to a number of challenges that can potentially jeopardize further sustainable operations, embracing advanced technologies becomes paramount. Thus, the ability to harness the power of ICT to deliver value will largely determine whether the energy industry can gain the upper hand in the digital transformation. Inspired by the beauty of modern technologies and the important role they are destined to play in the energy transition, the research presented in the current thesis aims to answer the question: **How to conduct a successful digitalization of energy systems?** Supported by the arguments presented in Chapter 1, the digitalization process is considered successful when it facilitates decarbonization and decentralization of the power sector while delivering long-term benefits for all actors along the energy value chain.

Taking into account how highly multifaceted the problem is, a specific research methodology that helps to address the power sector's digitalization is devised in Section 2.2. Thus, approaching the matter of digitalization in a highly practical manner, this thesis revolves around real-world implementations of digital technologies in energy with the following objectives underlying the research:

Objectives

1. Identify some of the most promising use cases of digitalization in the power sector, followed by developing and implementing ICT-based digital energy services.
2. Evaluate the impact of digitalization on achieving the energy transitions' overarching goals while determining the added value for utilities and their customers.
3. Provide recommendations for digitalizing the power sector, taking into account the course of development in energy systems and digital technologies.

2.2 Research methodology

The power sector's digitalization is a complex issue that covers the entire energy value chain and requires a holistic and structured approach to achieve the objectives defined in Section 2.1. This thesis investigates the digitalization phenomena from the energy management perspective, thus focusing on the ultimate link of the energy value chain, where emerging digital energy services are disrupting the established relationships between utilities and their customers. At this point, the power sector analysis should be conducted with the system mindset, taking into account all adjacent fields, including technology, economy, policy, and society. The scope of this thesis implies an in-depth analysis solely across technological dimension with further insights on economical and societal benefits. Thus, Figure 2.1 illustrates the proposed research framework to carry out an exhaustive analysis of energy systems within the digitalization phenomena.




Figure 2.1 – Proposed energy systems research framework

The energy process design, which forms the basis of any digital energy service, is represented by three consecutively connected layers. First, the ICT background is set in place that determines the key technological enablers for digitalizing any energy process. The construction of this layer often requires combining several digital technologies, thus a profound understanding of their operational principles, strengths, and weaknesses is required to enhance as opposed to complicate the process. Second, the appropriate use cases must be selected to define the context for reimagining the energy services and extracting the maximum value from digital technologies. Identifying such use cases is largely associated with determining the boundaries of the energy system, which, according to [Zhou 2016], is done best along the time, space, and user dimensions. The third layer specifies the objectives to be achieved by energy digitalization, where energy efficiency, environmental impact, cost-effectiveness, and energy security are commonly topping the list of desired outcomes. The following sections elaborate on applying the proposed research framework to investigate energy digitalization in the current thesis.

2.2.1 Selected ICT technologies

Considering the diversity of digital technologies presented in Section 1.1.3, selecting only some of them was necessary to narrow down the thesis's scope for conducting more in-depth research of the digitalization phenomena. The criteria underlying the selection process included the technology's maturity level, potential to significantly impact the electric power sector, expected breadth of application, and previous examples of successful implementation in other industries. The cost of introducing new technology and the complexity of respective organizational transformations were excluded from the selection criteria, as in the real-world implementation they would largely depend on the scale, resourcefulness, and structure of a specific energy utility. Following the aforementioned criteria, two digital technologies were selected to form the basis of digitalization research: the ML subset of AI and blockchain. Part I of this thesis tackles the application of both supervised and unsupervised ML algorithms, while Part II focuses on exploring the potential of RL and blockchain.

2.2.2 Energy use cases

The energy management use cases implicated in this study are characterized along three dimensions that determine the application scope of digital technologies. The time dimension defines the period of interest and the time resolution of utilized data. The majority of inputs used in this thesis come with 15-minute or hourly resolution, where the prior is common for measurements obtained from conventional smart meters, while the latter is often seen in weather data. The space dimension gives an idea of a system scale. Part I of the current thesis focuses on energy management at a building level, where both residential and tertiary sites are explored. Part II, instead, analyzes a group of tertiary buildings that form a smart community, although some of the results may deem applicable for individual households and residential communities. The last, user dimension, determines the beneficiary of a digitalized energy process, whether it is a single person in a household, a community of users in a country, or an organization, such as energy utility. This thesis deals mostly with owners and inhabitants of residential households and building managers and users of tertiary facilities.

2.2.3 Energy goals

Depending on specific digital technology and the respective application use case, the energy goals pursued in the current thesis vary. However, to produce meaningful results, the goals are endowed with common characteristics that render them measurable, realistic, attainable, and time-bounded. The overarching objective that this thesis aims to fulfill is the increase in energy efficiency through digitalization. Part I is centered around energy savings that stem from managing energy consumption in a user-centric manner and result in direct financial benefits. Part II focuses on optimizing energy processes to facilitate the integration of RES, improve energy security, reduce costs, and obtain better utilization of flexibility sources.

Energy Management in Smart Buildings

Part I

3 The Smart Building ICT Framework

Widespread digitalization promises to unlock buildings' massive potential for energy savings, endowing them with a layer of intelligence. However, efficient management of smart buildings' hardware and software assets is a non-trivial task that can be accomplished using emerging phenomena of platformization. This chapter introduces an ICT-based platform that facilitates the deployment of digital energy services and represents an interesting research framework to explore digital energy management solutions at the smart building level.

Related publications:

[3]^a Albert Hoffrichter, Evangelos Zacharis, Angelina Katsifaraki, Ashley Morton, Gloria Calleja, Federica Fuligni, Marko Batic, Marina Dorokhova, Niall Castelli, Konstantinos Kanellos and Thanh Nguyen. *Behavioral Change towards EE by Utilizing ICT Tools*. MDPI Proceedings of The 8th Annual International Sustainable Places Conference (SP2020), vol. 65, no. 1: 10, 2020. DOI: 10.3390/proceedings2020065010

[4]^b Filipe Soares, André Madureira, Andreu Pagès, António Barbosa, António Coelho, Fernando Cassola, Fernando Ribeiro, João Viana, José Ricardo Andrade, Marina Dorokhova, Nelson Morais, Nicolas Wyrsh, Trine Sørensen. *FEEDBACK: An ICT-based platform to increase energy efficiency through buildings' consumer engagement*. *Energies*, 14(6), 1524, 2021. DOI: 10.3390/en14061524

[5] Marina Dorokhova, Fernando Ribeiro, António Barbosa, João Viana, Filipe Soares, Nicolas Wyrsh. *Real-world implementation of an ICT-based platform to promote energy efficiency*. *Energies*, 14(9), 2416, 2021. DOI: 10.3390/en14092416

^aThe author contributed with the section describing FEEDBACK's technical and legal challenges.

^bThe author contributed to ICT-based platform's design and developed applications described in Chapters 4-6.

3.1 Background and motivation

The evolution of the energy value chain is determined by progressive changes happening to its key participants, assets, and processes. With consumers turning into prosumers and ubiquitous digitalization spreading throughout the power sector, the energy end-users strive to reinvent themselves to meet the needs and demands of the digital world. The buildings sector, responsible for 40% of final energy consumption in the EU, is the largest energy consumer followed by transportation and industry [EC 2021]. Therefore, buildings' digitalization represents a massive potential for infrastructure owners, operators, and facility managers to reduce energy use, minimize environmental impact, and improve the well-being of buildings' occupants. The buildings of the digital era are called smart buildings, with the collective definition given as follows:

Smart Building

A structure enhanced with sensors, actuators, telecommunication links, and computer-based algorithms, capable of sensing, monitoring, and automatically controlling its operations.

The complex concept of smart buildings requires new managerial instruments to administer and control the expanding hardware and software assets successfully. Going beyond simple automation, the new-generation of Building Management Systems (BMSs) is supported by the emerging platformization phenomena in the energy field [Kloppenburger 2019]. The ICT-based platforms consist of diverse interconnected single-purpose applications that provide digital energy services operating on top of existing infrastructure. Endowing buildings with a layer of intelligence, these platforms constitute interesting research frameworks to explore the digital energy management solutions at the smart building level. The research conducted in Part I of the current thesis evolves around such ICT-based platform developed in the course of the FEEdBACK project.

3.2 The FEEdBACK project

The FEEdBACK project funded by the European Commission under the H2020-EE-2016-2017 call stands for Fostering Energy Efficiency and Behavioral Change through ICT. The FEEdBACK project envisages the development of an innovative ICT-based platform to promote energy-efficient behaviors among buildings' users and to aid building managers in maximizing energy savings while achieving other advantages, such as peak load reduction and increase of local RES consumption. The ICT-based platform is expected to unite various software applications and a gamification engine into the interactive energy management system that aids interested stakeholders to optimize energy consumption, generation, and storage of controllable and flexible devices without compromising comfort levels. The specific objectives of the FEEdBACK project include the following:

- Exploit innovative and user-friendly ICT solutions to promote users' and interested stakeholders' engagement and to foster energy efficiency
- Explore intelligent control and automation to optimize the utilization of controllable resources
- Include insights from social and behavioural sciences to maximize permanent behavioural change
- Maintain/improve comfort levels and indoor air quality
- Make energy usage data accessible to users and to designated third parties
- Setup demonstrators with different characteristics and quantify impacts using rigorous baselines
- Investigate the potential for replication to other boundary conditions

The project's consortium is composed of 8 partners based in 7 different countries – Portugal, the Netherlands, Switzerland, Spain, the United Kingdom, Denmark, and Germany. The demonstration sites participating in the project include 3 different European regions: Porto (Portugal), Barcelona (Spain), and Lippe (Germany). The Porto test site is located at the INESC TEC research institute's headquarters and consists of two contiguous building sections that cumulatively host about 400 occupants. Despite being built in different years, the two buildings are similar in their external and internal structure. The Barcelona demonstration site is situated in the town of El Prat de Llobregat and encompasses 10 various tertiary buildings owned by the municipality. The buildings include offices, sports centers, cultural spaces, and educational facilities. The Lippe test site is located in the village of Dörentrup and is represented by more than 30 private residential households.

3.3 The ICT-based smart building platform

Figure 3.1 depicts the architecture of the FEEDBACK ICT-based platform, where key components can be divided into 4 large groups: external inputs, DEXCell, gamification platform, and applications.

External inputs consist of meters and sensors installed in demonstration sites. The smart meters measure energy consumption and multi-sensors provide information about temperature, humidity, illuminance, and equivalent CO₂ concentration, thus describing the conditions of the ambient environment.

DEXCell is a tool provided by a FEEDBACK partner that encompasses data storage capabilities with Graphical User Interface (GUI) referred to as FEEDBACK Suite. Thus, most data fed or produced by the ICT-based platform is stored in DEXCell and visualized through GUI.

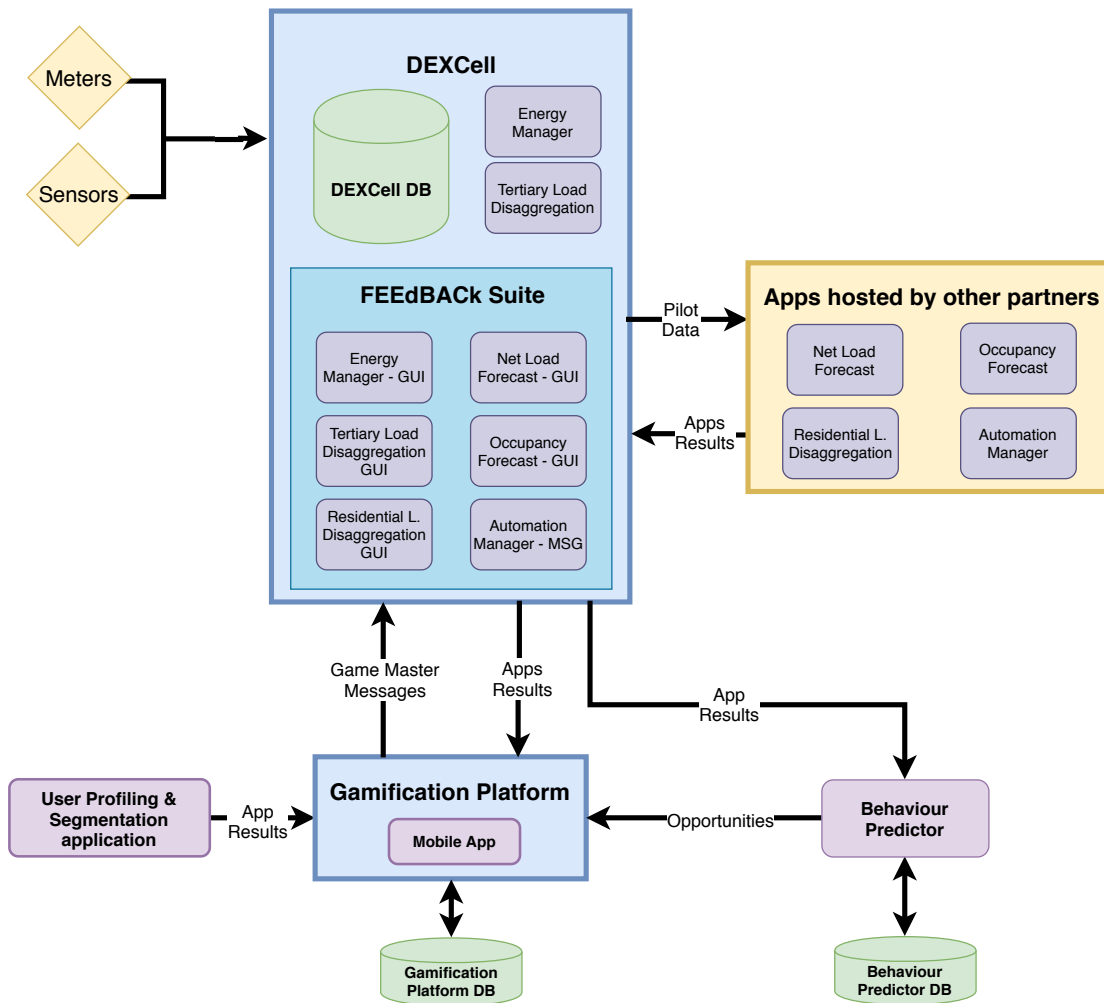


Figure 3.1 – The architecture of the FEEdBACK ICT-based platform

Gamification platform acts as the interface between the ICT-based platform and the end-user aiming to encourage more efficient energy utilization and more responsible consumer behavior by providing the user with an interactive experience through a mobile application. The latter represents a user-friendly GUI that displays individual applications' results, assists the buildings' occupants and managers with decision-making, and fosters users' engagement through personalized and opportunistic messages and gamified peer competition.

Applications developed by different project partners represent interconnected digital energy services that, due to their specificity, are hosted on partners' own servers. Notably, the energy manager and tertiary load disaggregation applications are hosted in DEXCell as they were developed by the same FEEdBACK partner that provided this tool. Other applications include behavior predictor, user profiling and segmentation, net load forecasting, residential load disaggregation, occupancy forecasting, and automation manager. The last 3 applications were developed by EPFL and are thus detailed further in the current thesis.

- The **Behavior predictor** app resorts to ML techniques and serves the purpose of identifying opportunities for lowering electricity consumption and assessing their impact on users' behavioral change and realized energy savings.
- The **User profiling and segmentation** app aims to characterize the users by attributing them to specific user profiles determined through an extensive questionnaire that considers motivational drivers and barriers for changing energy behavior, pro-environmental beliefs, values, and attitudes.
- The **Net load forecasting** app consists of several modules that provide the possibility to forecast PV generation, building's baseline energy consumption, and specific load demand of various appliances. The application uses supervised ML algorithms to build reliable short-term forecasts based on historical on-site data and exogenous meteorological variables.
- The **Energy manager** app combines the inputs from the gamification engine and the platform's back-end to provide decision-making support to facility managers. Thus, the app displays short messages that help to identify energy-saving opportunities within buildings and respective controlling actions to be taken.
- The **Load disaggregation** app consists of tertiary and residential modules that aim to provide an estimated disaggregated energy consumption from whole-site smart meter measurements in a non-intrusive manner. Both modules utilize historical energy consumption and other exogenous variables to transform the aggregated smart meter recordings to specific energy consumption values allocated to categories of devices. The application's output is used to raise the awareness of energy consumption behaviors among building users and identify energy-saving opportunities. The residential load disaggregation digital service is depicted in Chapter 4 of the current thesis.
- The **Occupancy forecasting** app supports other applications by indicating the building's occupancy for the day ahead. Depending on the demonstration site, the app provides forecasted presence and absence information by sub-metering zone, by functional area, or by a house as a whole. The algorithms at the core of the application are context-based, thus the provided occupancy information depends entirely on sensing the conditions of indoor climate or approximating it through electricity usage. Chapter 5 of this thesis presents the details of the occupancy forecasting digital energy service.
- The **Automation manager** app aims at producing day-ahead energy-efficient schedules to operate controllable loads, such as Heating, Ventilation, and Air Conditioning (HVAC), without compromising the users' comfort. The generated optimized schedules are supplied to building managers and household owners for verification and deployment. The latter resorts to automated load control enabled by BMS present in demonstration sites. Chapter 6 of the current thesis elaborates on the automation digital energy service.

4 Load Disaggregation

To efficiently respond to the challenge of balancing demand and supply in the context of increasing penetration of renewables, demand-side management measures are implemented at the consumer's side. However, their success depends largely on the availability of detailed knowledge of one's energy consumption patterns. To understand what drives the energy usage in residential households and raise the awareness of one's energy-related behavior without compromising privacy and security, a non-intrusive load monitoring method is implemented.

The main **highlights** and **contributions** of this chapter are:

- The introduction of a novel unsupervised load disaggregation algorithm based on activity-chain reconstruction through the Markov process. The low-sampling-rate aggregated load curves are disaggregated into 8 distinct categories of household appliances.
- The algorithm is benchmarked to four state-of-the-art load disaggregation algorithms on three publicly available datasets. The results show that the algorithm achieves similar disaggregation accuracy while exhibiting lower computational cost.
- The algorithm's detailed performance evaluation on real-world datasets from Norway and Germany demonstrates that obtained prediction uncertainties do not exceed 20%. The results highlight the importance of considering regional characteristics, behavioral patterns, and consumer preferences to provide a better disaggregation experience.

Related **publications**:

[1]^c Jordan Holweger, Marina Dorokhova, Lionel Bloch, Christophe Ballif and Nicolas Wyrsh. *Unsupervised algorithm for disaggregating low-sampling-rate electricity consumption of households*. Sustainable Energy, Grids and Networks, vol. 19, page 100244, 2019. DOI: 10.1016/j.segan.2019.100244

^cThe author contributed to the literature review and benchmarking of the algorithm's performance.

4.1 Background and motivation

Transitioning to a digitalized energy landscape while facing new challenges, such as increasing penetration of RES, augmenting self-consumption, and balancing demand and supply, requires new approaches to realize energy management duties successfully. Demand-Side Management (DSM) refers to a group of measures designed to efficiently manage the energy consumption at the consumer's side while providing utilities with the opportunity to reduce incurred costs and improve the electrical grid's reliability and efficiency [Palensky 2011]. Secured by financial incentives and behavioral change through energy awareness, implementation of DSM requires the availability of new-generation smart devices and a profound understanding of users' energy consumption patterns. While the prior is actively supported by the proliferation of smart meters and the advent of smart controllable appliances, the latter depends largely on the availability of relevant data and skillful implementation of AI-enabled techniques.

From the consumer's perspective, possessing a detailed knowledge about one's energy consumption helps greatly make informed decisions regarding achieving a more sustainable energy behavior, realizing potential energy savings, and purchasing new electrical appliances. Implementing specific digital energy services can help consumers accomplish the aforementioned goals while providing a personalized and interactive experience. Thus, monitoring, controlling, and conserving energy in buildings can be obtained by realizing a four-step procedure embedded in the ICT-based platform. First, the energy consumption must be measured with adequate time resolution. Second, the data must be analyzed to identify prominent consumption patterns, determine energy-saving opportunities, and provide guidelines for modifying energy-related behavior. Third, real actions must be taken to tap into identified opportunities. Last, the evolution of energy consumption must be tracked to determine the progress in carrying out the energy-saving efforts. However, the implementation of such a procedure poses a greater challenge to maintain consumer's privacy and security. Thus, the unconventional methods to distill the energy consumption data into actionable energy insights, such as Non-Intrusive Load Monitoring (NILM), have to be implemented.

The remainder of this chapter is organized as follows. Section 4.2 summarizes state of the art in the NILM literature, particularly focusing on various ways to categorize different studies. Section 4.3 presents the methodology underlying the proposed low-frequency unsupervised algorithm and discusses the benchmarking procedure conducted against other popular NILM techniques. Section 4.4 concentrates on the suggested algorithm's applicability and evaluates its performance on two case studies according to designated metrics. Finally, Section 4.5 concludes the findings of this study and highlights several key features that would greatly improve the adoption of such digital energy service in the future.

4.2 State of the art

The NILM, introduced by Hart in 1992 [Hart 1992], refers to the method of decomposing the aggregated load curves into power signals of individual appliances without explicitly measuring them. As various approaches to disaggregation emerged, the need to classify the developments in the field of NILM became obvious. Thus, the division of NILM algorithms into several categories, such as high and low frequency, supervised and unsupervised, residential and industrial, was proposed.

The division by frequency suggested in [Esa 2016], groups the NILM approaches into low-frequency and high-frequency categories, where the prior corresponds to sampling rates below 1 Hz, and the latter typically ranges between 1 kHz and 0.5 MHz. Despite being promising, the high-frequency methods based on transient analysis of power signals demonstrate low cost-effectiveness due to limited sensing capabilities of conventional smart meters and high requirements for data transmission [Wang 2018]. Instead, the low-frequency methods utilize the smart meter power measurements obtained with sampling intervals of 5, 15 minutes, or longer. The examples of deploying NILM techniques at extremely low frequencies include the works of [Zhao 2018, Kolter 2010] and [Batra 2016]. The prior tackled the disaggregation problem at the hourly resolution while the latter extracted categorized end-use energy consumption from monthly electricity bills.

Another way to categorize disaggregation algorithms is to split them into supervised and unsupervised as suggested in [Faustine 2017]. The supervised NILM algorithms manage to recognize individual appliances' power signals in the aggregated load curves due to extensive preceding training on labeled data. Although the unsupervised methods do not require such preliminary treatment and can perform disaggregation duties directly, their implementation comes at the expense of greater uncertainties in disaggregation results. The supervised algorithms successfully deploy classic ML techniques, such as ANNs and decision trees, which were demonstrated in [Kelly 2015a, Biansoongnern 2016] and [Liao 2014], respectively. Another popular supervised method, presented in [Kumar 2016, Kumar 2017], utilizes Graph Signal Processing (GSP) to investigate the underlying graph structure of the aggregated power signal. The unsupervised NILM algorithms commonly deploy the Factorial Hidden Markov Model (FHMM), as discussed in [Bonfigli 2015] and implemented in [Zoha 2013].

Both methods, supervised and unsupervised, despite their previous successes, come with drawbacks that hinder their implementation in the real world. The deficiency of supervised NILM algorithms is indeed the requirement to utilize labeled data, which are often not public or generalizable. Thus, such methods are deemed suitable only for general studies as their application to actionable DSM is limited. The existing unsupervised algorithms are typically operated at relatively high sampling frequencies making the necessary data collection costly and difficult. Moreover, they rarely allow a detailed disaggregation of the load curve into categories of devices by rather focusing on identifying the consumption patterns that stand out. Therefore, a need for a low-sampling rate unsupervised disaggregation algorithm emerged.

4.3 Methodology

4.3.1 Device Usage Estimation algorithm

The Device Usage Estimation (DUE) algorithm, presented in Figure 4.1, is a hybrid method that lies in-between the load simulation and load disaggregation. Operating at a 15-minute sampling rate, the DUE algorithm aims to disaggregate the whole-house active power measurements into categories of appliances rather than infer their exact power signals. The total of 8 categories, namely cooking, entertainment, fridge, heating, housekeeping, ICT, light, and standby, is defined. The inputs required for the algorithm include the household's daily load curves L obtained from the smart meter, the day type D , characteristics of inhabitants, and the list of appliances in the household's possession. The age-group G , employment status of inhabitants E , and the household's inventory data can be conveniently collected through the digital survey during the household's registration on the ICT-based platform.

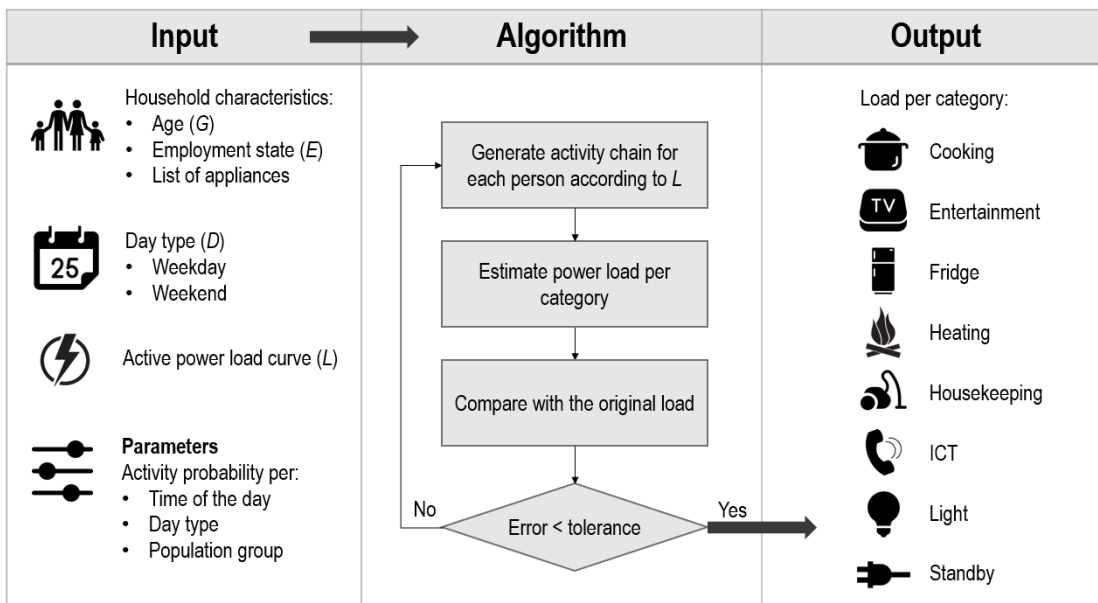


Figure 4.1 – The Device Usage Estimation (DUE) algorithm explained

The DUE's methodology uses a Markov process to generate daily activity chains for each person in the household. The switch between activities during the course of the day is characterized by the transition matrix, while each individual activity can be described with its initial and duration probabilities. The coefficients of the transition matrix and activities' initial probabilities that depend on D, E, G were extracted from the time-of-use survey conducted in the Netherlands [Sociaal en Cultureel Planbureau 2005]. Although such geographical choice is assumed to be representative for determining European behavioral patterns, specific surveys of a similar nature can be organized through an add-on application on the ICT-based platform. Once the activity chain is generated, the probable list of appliances used throughout the day can be inferred based on relations described in Table 4.1.

Table 4.1 – List of possible activities and related appliances

Activity	Appliances
Cleaning	vacuum, TV, stereo, lights
Using a computer	TV, stereo, PC, laptop, printer, lights
Cooking	stove, oven, microwave, kettle, TV, stereo, lights
Washing dishes	dishwasher, TV, stereo, light
Eating	coffee maker, microwave, kettle, TV, stereo, lights
Doing the homework	TV, stereo, PC, printer, laptop, lights
Playing a game	TV, stereo, gaming console, lights
Laundry	washing machine, tumble dryer, TV, stereo, lights
Music	stereo, PC, tablet, laptop, lights
Watching TV	TV, DVD player, PC, tablet, laptop, lights
Showering	hairdryer, TV, stereo, lights
Outdoor, sleeping, working	∅

As each activity is characterized by possible simultaneous usage of multiple devices, their cumulative power demand, determined as the sum of individual appliances' nominal powers, should respect the available power budget imposed by the household's load curve L . Therefore, for each identified activity, the DUE algorithm strives to find a combination of appliances present in the household's inventory that respects the available power budget with specific predefined tolerance. However, estimation of standby, fridge, and light differs from other categories. Standby that is filtered out first corresponds to the minimum value of the load curve at consideration, so it also captures a constant portion of the fridge's consumption. The fridge's typical power signal is extracted using the clustering of nightly load curves. Once the pattern is identified, the fridge cycle's start and duration are adjusted daily to fit the observed load curve best. Finally, the light consumption accompanying most activities listed in Table 4.1 is calculated between sunset and sunrise according to active occupancy instances determined based on detected power peaks. Once appliances are allocated to detected activities in the generated activity chain, their simulated power signals are aggregated to form 8 disaggregation categories, as shown in Figure 4.2.




Figure 4.2 – Example of a load curve disaggregated using DUE algorithm

Devices' nominal powers, extracted from [Kelly 2018], and their affiliation to predefined categories are presented in Table 4.2. Notably, despite mentioning several appliances, the DUE algorithm's current version does not deal with space heating. Thus, only the hairdryer is included in the heating category. As technology progresses and the market of household devices evolves, the need to periodically update the list of household appliances emerges. In the digitalized world, such a simple action can be conveniently handled by the ICT-based platform, where smart building's inhabitants would not only manage their household's inventory online but receive a timely reminder to do so. Hence, feeding the ICT-based platform with an accurate description of the devices owned contributes to the performance of the NILM algorithm and provides a truly personalized experience.

Table 4.2 – Appliances and their nominal power grouped per category

Category	Appliance	P_{Nominal} [W]
Cooking	coffee maker	800
	microwave	1250
	kettle	1800
	oven	2400
	stove	500
Entertainment	TV	124
	TV box	20
	DVD player	80
	PC	110
	laptop	55
	tablet	7
	stereo	100
	gaming console	180
Fridge	fridge (with a freezer)	94
	fridge (without a freezer)	66
	freezer alone	62
Heating	hairdryer	600
	boiler	2000
	heat pump	1000
Housekeeping	washing machine	406
	tumble dryer	2500
	dishwasher	1131
	vacuum	2000
ICT	printer	23
Light	lighting	137
Standby	modem (and similar)	8

4.3.2 Benchmarking

To assess the validity and performance of the unsupervised DUE method, it was benchmarked against the following algorithms, as detailed in [Holweger 2019]:

- Combinatorial Optimization (CO) [Hart 1992]
- Factorial Hidden Markov Model (FHMM) [Kim 2011]
- Graph Signal Processing (GSP) [Kumar 2017, He 2016]
- Discriminative Disaggregation via Sparse Coding (DDSC) [Kolter 2010]

The benchmarking procedure was carried out on three publicly available datasets from various regions in Europe: ECO [Beckel 2014], SMARTENERGY.COM [Alhamoud 2014], and UK-DALE [Kelly 2015b]. The power signals of individual devices, contained in each dataset, were down-sampled to a 15-minute resolution and further aggregated to form the appliances' categories. The duration of the available training and testing data is presented in Table 4.3 for each of the datasets.

Table 4.3 – Testing and training periods for the three datasets

	Training period			Testing period		
	from	to	days	from	to	days
ECO	2012-06-01	2012-09-30	121	2012-10-01	2012-10-30	29
SMARTENERGY.COM	2013-03-05	2013-05-19	75	2013-05-19	2013-06-25	37
UK-DALE	2014-04-01	2015-04-01	365	2015-04-01	2015-07-01	91

To measure the performance of the DUE algorithm and compare it quantitatively with chosen benchmark methods, the following metrics adapted from [Makonin 2015] were utilized:

$$\text{Estimation Accuracy per category } m: \text{ESTACC}_m = 1 - \frac{\sum_t |\hat{P}_m^t - P_m^t|}{2 \cdot \sum_t P_m^t} \quad (4.1)$$

$$\text{Estimation Accuracy: ESTACC} = 1 - \frac{\sum_t \sum_m |\hat{P}_m^t - P_m^t|}{2 \cdot \sum_t \sum_m P_m^t} \quad (4.2)$$

Where P_m^t and \hat{P}_m^t are the ground truth and the estimated power signal of category m , defined in Table 4.2, respectively. The additional *Energy Share Error* metric, detailed in Equation 4.3, is used to evaluate the quality of each category's disaggregation with respect to the total energy consumed. Thus, the energy share of a particular category m is calculated for both the ground truth and the estimation cases, and the error is determined as the difference between the two.

$$\text{Energy Share Error per category } m: \text{ESE}_m = \frac{\sum_t \hat{P}_m^t}{\sum_t \sum_m \hat{P}_m^t} - \frac{\sum_t P_m^t}{\sum_t \sum_m P_m^t} \quad (4.3)$$

Figure 4.3 depicts the *Energy Share Error* for each algorithm across all datasets. Although this metric depends largely on both the dataset and the categories observed, the median values of *Energy Share Error* per category do not exceed 20%. The DUE algorithm's performance is slightly inferior to benchmark algorithms' performance due to its unsupervised nature. However, one should consider that contrary to other algorithms the DUE algorithm always assumes the presence of all devices' categories within the household. Thus, the missing housekeeping category in the SMARTENERGY.COM dataset and the absence of heating in the ECO dataset contribute to the discrepancy between algorithms' results.

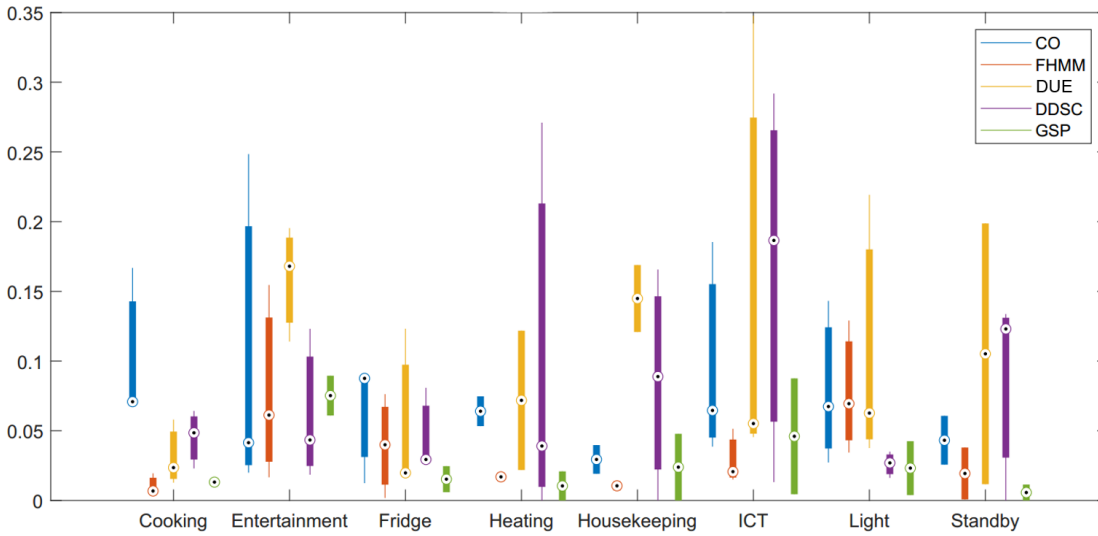


Figure 4.3 – *Energy Share Error* for all algorithms across all datasets

Figure 4.4 presents the *Estimation Accuracy* achieved by algorithms on each test dataset and their respective execution time. Despite the mediocre accuracy of the DUE algorithm, its sequential nature allows it to scale linearly with the dataset's length and results in competitive computational speed. Among the *Estimation Accuracy* and the *Energy Share* metrics, the latter, due to its non-temporal nature, is considered to be more comprehensive and user-friendly when communicating the disaggregation results to the users of the ICT-based platform.

To conclude, during the extensive benchmarking procedure, the DUE algorithm has proved its applicability to solve the NILM problem and demonstrated comparable performance to supervised disaggregation algorithms. Moreover, it has shown that its key value for both utilities and household's inhabitants lies in its unsupervised nature. With the prior often missing the reference energy consumption to perform the training and the latter being generally concerned with preserving their privacy, the unsupervised DUE algorithm fits both by eliminating the need for detailed metering from within the household. The following sections will reflect on the real-world implementation of the DUE algorithm and its suitability for digital energy services.

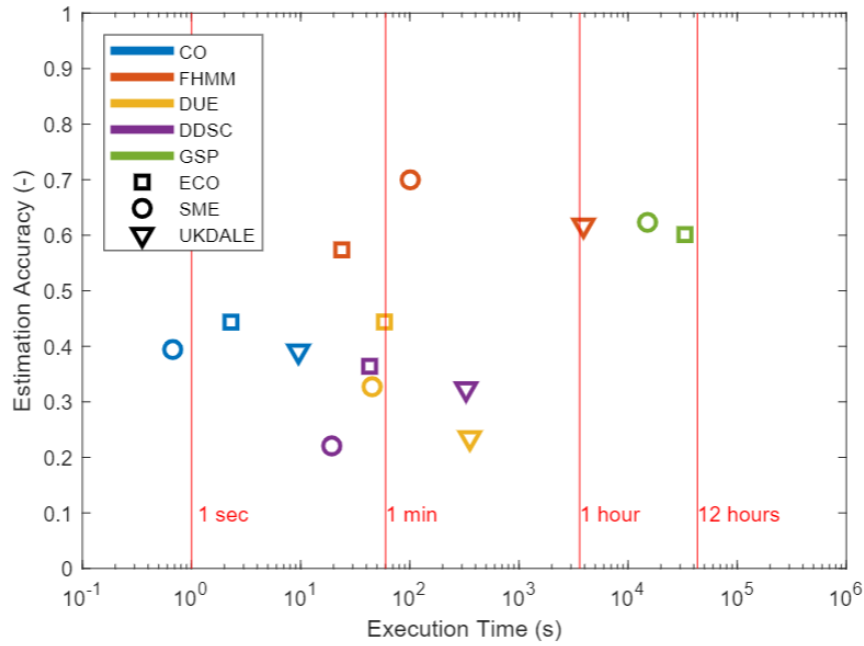


Figure 4.4 – *Estimation Accuracy* versus execution time for all algorithms across three datasets

4.4 Case studies

4.4.1 Norway

The Norwegian case study occurs in two households, an independent family house and a couple apartment, during the test period from 2017-12-17 to 2018-01-19. The data, essential for DUE's input, such as characteristics of the households and detailed lists of domestic appliances, were obtained through a digital questionnaire and are presented in Table A.1.

Table 4.4 summarizes the main appliances and their respective categories present in both households' ground truth. One can notice that despite several disaggregation categories existing in the house and apartment labeled data, for most of them, only a few devices are declared in the ground truth. Therefore, such under-representation of sub-metered appliances compared to households' inventories is likely to cause overestimation of the energy usage per category. Another particularity of collected labeled data is related to the house's heating category represented by the heat pump. As disaggregation of a heat pump is not yet implemented in the current version of the DUE algorithm, the heat pump's power measurements were subtracted from the house's aggregated smart meter data before the disaggregation process. The apartment, instead, is connected to the district heating network.

Table 4.5 presents the load disaggregation results obtained in Norwegian demonstration households using the DUE algorithm. For both households, the *Energy Share Error* per category is calculated based on the ground truth and disaggregation outcome. Categories absent in the ground truth are marked with n/a.

Table 4.4 – Appliances present in the ground truth of Norwegian demonstration households

Category	House	Apartment
Cooking	coffee maker	kettle oven
Entertainment	TV	TV
Fridge	cooling unit	n/a
Heating	heat pump	n/a
Housekeeping	dishwasher washing machine tumble dryer	dishwasher tumble dryer

Table 4.5 – Load disaggregation results of Norwegian households: ground truth and predicted *Energy Share* [%] per category followed by the *Energy Share Error* [%] per category

Categories	House			Apartment		
	Ground truth	Predicted	ESE	Ground truth	Predicted	ESE
Cooking	0.1	8.0	7.9	1.7	10.9	9.2
Entertainment	0.1	1.8	1.7	5.1	7.2	2.1
Fridge	0.3	13.9	13.6	n/a	19.7	n/a
Heating	0.0	0.1	0.1	n/a	0.4	n/a
Housekeeping	0.4	6.2	5.8	7.5	8.1	0.6
ICT	n/a	0.4	n/a	n/a	0.5	n/a
Light	n/a	4.8	n/a	n/a	5.6	n/a
Standby	n/a	64.8	n/a	n/a	47.6	n/a

Following the DUE's validation results presented in Figure 4.3, the algorithm's application to Norwegian demonstration households exhibits the *Energy Share Error* per category within the 15% range. In both cases, the disaggregation algorithm expectedly overestimates the energy usage per category as more appliances are present in the households' inventories than declared in the ground truth. Such phenomena is particularly noticeable in the cooking and entertainment categories. The house's fridge category is largely overestimated as, in addition to the cooling unit present in the ground truth, the household possesses a fridge and freezer.

To compensate for the low number of household appliances comprising the ground truth, the quality of disaggregation results is additionally evaluated on statistics of Norwegian residential electricity end-use depicted in Figure 4.5. As the DUE algorithm's current version does not include space heating disaggregation, the energy consumption breakdown is presented with and without heating to account for substantial differences in the set of devices attributed to this category. Being the largest energy consumer in Norway, heating is especially fundamental during the winter season. More than 70% of Norwegian households typically use a combi-

nation of various energy sources to heat their homes, with electricity and firewood being the most common [Dalen 2013]. Moreover, the space heating appliances diversify largely according to their working principle. Beyond simple resistance heaters, they include other heating equipment that cannot be characterized with a fixed amount of rated power. Therefore, the complexity of disaggregating some of the most typical heating options in Norway, such as heating cables, panel convective heaters, infrared heaters, floor heating, and sauna, contributes to disaggregation errors and uncommonly high standby consumption.




Figure 4.5 – Breakdown of residential end-use electricity demand in Norway 2006 [Dalen 2013]

As seen in Table 4.6, the *Energy Share Error* of the fridge, housekeeping, and light categories before excluding heating is relatively small, thus indicating that predicted energy shares fall in the ballpark of statistical energy usage shown in Figure 4.5. The same metric for the ICT category demonstrates larger underestimation as only one device, printer, is considered by the algorithm, whereas in reality, more appliances might belong to this category. Removing heating increases the *Energy Share Error* of housekeeping, ICT, and light while maintaining the values within the 10% range. Although the fridge category's behavior differs for house and apartment, the metric equally does not exceed 10%. The *Energy Share Error* of other category that encompasses cooking, entertainment, and standby decreases significantly after heating removal due to better accommodation of the unusually high energy share of standby inherent to both house and apartment.

Chapter 4. Load Disaggregation

Table 4.6 – *Energy Share Error* [%] per category deduced from comparing load disaggregation results to Norwegian statistics of electricity end-use with and without heating

Categories	House		Apartment	
	with heating	without heating	with heating	without heating
Fridge	0.7	-6.5	6.5	-0.7
Housekeeping	-2.6	-7.5	-0.7	-5.6
ICT	-5.4	-8.6	-5.3	-8.5
Light	0.2	-2.3	1.0	-1.5
Other	42.4	24.8	33.5	15.9

Interestingly, the standby category turned out to be the largest consumer in both demonstration households, although the apartment's inhabitants declared no special heating arrangements. Therefore, either the imperfection of the chosen standby detection methodology or the oddity of the Norwegian energy consumption pattern can be suspected. Figure 4.6 helps to clarify the problem with standby, as one can notice that disaggregated categories do not fit perfectly the observed load curve depicted in green. Therefore, the energy share of standby gets overestimated, as the total load resulting from the sum of all categories' loads appears to be smaller than the observed load.




Figure 4.6 – Load disaggregation example of one day in Norwegian apartment

4.4.2 Germany

The German case study investigates 10 private residential households during the period of 7 months lasting from 2019-08-01 to 2020-02-29. Each household's characteristics and respective lists of appliances owned were collected in the framework of the FEEDBACK project and are presented in Table A.2. The aggregated power measurements were obtained from households' smart meters with 15-minute resolution. However, the individual appliances' ground truth was not collected due to the absence of sub-metering inside the households. To preserve the inhabitants' privacy, each house is associated with a unique identification number. All chosen households can be divided into two equal groups according to the number of occupants: inferior or equal to 2 and greater than 2.

Table 4.7 presents the load disaggregation results in German households, obtained using the DUE algorithm. Similarly to the Norwegian case, one can notice the large shares of standby consumption for certain households. As discussed previously, the lack of proper heating disaggregation together with poor peak load filling contribute to overestimating standby. To demonstrate the influence of unfilled load peaks during the disaggregation process, the expected *Energy Share* of standby was calculated in addition to the predicted one:

$$\text{Standby power: } \hat{P}_{standby}^t = \min_t(P^t) \forall t \in T \quad (4.4)$$

$$\text{Expected standby Energy Share: } ES_{standby}^{expected} = \frac{\sum_t \hat{P}_{standby}^t}{\sum_t P^t} \quad (4.5)$$

$$\text{Predicted standby Energy Share: } ES_{standby}^{predicted} = \frac{\sum_t \hat{P}_{standby}^t}{\sum_t \sum_m \hat{P}_m^t} \quad (4.6)$$

Where P^t is the aggregated power signal of a household at time t , T is the time domain of the considered day, and $\hat{P}_{standby}^t$ defined in Equation 4.4 corresponds to the detected standby power. The difference between the two ways of computing the standby's *Energy Share* lies in the total energy consumed during the test period. Particularly, the expected *Energy Share* is determined based on the observed aggregated power signal P^t as stated in Equation 4.5. Instead, the predicted *Energy Share* imposed by Equation 4.6 is calculated relative to the power signal $\sum_m \hat{P}_m^t$ reconstructed for all categories m . As can be seen in Figure 4.7, unfilled load peaks result in total disaggregated load $\sum_m \hat{P}_m^t$ being inferior to the observed load curve P^t . Consequently, the denominator of Equation 4.6 appears to be smaller than the denominator of Equation 4.5, thus the predicted standby's *Energy Share* $ES_{standby}^{predicted}$ exceeds the $ES_{standby}^{expected}$.

Table 4.7 – Load disaggregation results for German households: predicted *Energy Share* [%]

	ID1002	ID1004	ID1007	ID1009	ID1011	ID1012	ID1013	ID1014	ID1017	ID1022
Cooking	5.3	3.0	4.5	10.9	10.6	7.0	11.3	6.5	5.3	13.1
Entertainment	0.7	0.9	1.7	0.7	2.0	1.7	1.4	0.9	1.7	0.7
Fridge	11.7	27.0	29.0	45.5	23.4	47.4	39.5	55.0	32.3	23.5
Heating	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Housekeeping	1.4	2.2	9.8	1.6	4.6	2.1	8.5	19.1	11.7	6.0
ICT	0.1	0.9	1.9	0.4	0.8	0.0	0.4	0.0	0.0	0.1
Light	6.2	5.8	6.9	2.5	6.2	5.8	3.9	5.1	11.1	7.7
Standby	74.6	60.2	46.2	38.4	52.4	36.1	31.1	13.4	37.9	48.9
Standby (real)	56.1	57.1	38.4	34.3	32.0	26.4	17.2	12.0	29.9	27.9




Figure 4.7 – Load disaggregation example of one day in German household ID1007

With the ground truth lacking from the households' observations, the statistical breakdown of electrical energy consumption in Germany's private households [Schweiss 2006], shown in Figure 4.8, is used to evaluate the disaggregation results. Excluding the standby discussed previously, the median *Energy Share Error* per category, summarized in Table 4.8, is within the 20% range. For most of the categories, the DUE algorithm has underestimated the *Energy Share*, with cooking being the closest to statistical energy consumption and housekeeping being the farthest. The fact that space heating disaggregation is not implemented in the current version of the DUE algorithm can explain the overestimation of the fridge category. According to the DUE's methodology, the fridge's energy consumption is inferred in the night hours by searching for symmetrical and repetitive peaks. Thus, some of the load peaks related to heating could be misinterpreted as the fridge load pattern.




Figure 4.8 – Breakdown of electrical energy consumption in Germany's private households 2006 [Schweiss 2006]

Table 4.8 – German demonstration households' median *Energy Share Error* across categories

Heating	Fridge	Housekeeping	ICT	Light	Cooking	Entertainment
-11.5%	14.8%	-15.3%	-11.9%	-5.1%	-1.6%	-9.9%

Chapter 4. Load Disaggregation

However, one should keep in mind that statistical electrical energy consumption presented in Figure 4.8 dates back to 2006, while the disaggregation test dataset was collected in late 2019. Therefore, the technological progress of various domestic appliances and increased energy consumption efficiency could have contributed to the discrepancy of *Energy Share* partitioning between categories. To enrich the analysis of the German households' load disaggregation, the comparison with the modern labeled GELAP dataset [Wilhelm 2021], collected in 20 private houses in Germany during the same test period, is provided in Table 4.9. As only a few sub-metered devices characterize each household in the GELAP dataset, the houses are aggregated into two groups according to the number of inhabitants. A further distinction is made based on the day type, resulting in weekday and weekend disaggregation.

Table 4.9 – Categories' *Energy Share* [%] comparison between GELAP dataset [Wilhelm 2021] and German demonstration dataset. The data is grouped according to the number of household's inhabitants and weekday/weekend feature.

Categories	GELAP dataset				German dataset			
	House ≤ 2		House > 2		House ≤ 2		House > 2	
	w-day	w-end	w-day	w-end	w-day	w-end	w-day	w-end
Cooking	3.2	3.3	4.9	6.5	9.8	9.7	5.6	3.9
Entertainment	9.8	6.9	5.9	5.9	1.0	1.0	1.2	1.6
Fridge	1.2	n/a	n/a	n/a	25.6	25.8	37.1	37.6
Heating	0.4	0.4	0.4	0.4	0.0	0.0	0.0	0.0
Housekeeping	3.2	3.3	4.9	6.5	3.5	2.8	9.1	8.1
ICT	2.0	1.5	0.6	0.8	0.1	0.1	0.5	0.5
Light	1.4	0.8	0.2	0.2	5.6	5.4	6.6	7.1
Other	78.8	83.8	83.1	79.7	54.3	54.9	39.8	41.1

Despite the GELAP dataset providing labeled appliance-specific power measurements, most households' energy consumption is not sufficiently explained, as the other category reaches 81.3% on average across all documented households. Moreover, the other category's large share can be characterized by the inclusion of domestic devices, which cannot be attributed to other categories, such as automatic garage door openers, circulation and water pumps. Despite the fridge category suffering from under-representation in the GELAP dataset, its disaggregation results together with the heating category can be explained by the disaggregation inaccuracies determined previously. For both datasets, housekeeping usage increases with the number of inhabitants in the household as supposedly the cleaning is needed more often. However, the weekday and weekend partitioning of housekeeping duties can be explained by occupants' preferences. Simultaneously, the magnitude of the energy attributed to this and light category varies potentially with the size of the household. The results of cooking, ICT, and entertainment differ between datasets, despite being in the same ballpark. Although the DUE algorithm is based on the statistical time-of-use survey, people's real habits and preferences related to daily activities might not fit the statistical profile, thus exhibiting a stochastic nature.

4.5 Conclusion

To conclude, this chapter presented the novel unsupervised load disaggregation method, the DUE algorithm, and discussed its application to Norwegian and German case studies. The suggested NILM methodology can serve as the basis for digital energy service that encompasses the detailed knowledge of one's energy consumption together with tailored recommendations for modifying energy-related behavior. The conducted case studies have highlighted the three key features that would greatly improve the accuracy and adoption success of such digital energy service in the future:

- To provide insights that matter and establish strong relationships with customers, the digital NILM service should enable greater personalization. The difference in disaggregation results between German and Norwegian households has clearly shown that regional characteristics, variations in household equipment, especially related to heating, behavioral patterns, and consumer preferences have to be considered. Thus, the usage of country-specific time-of-use surveys and expansion of the appliances regarded in the disaggregation process are paramount.
- To unlock the efficiency in load management and rapidly spot the inaccuracies in disaggregation, the customers' inventories of household appliances should be regularly updated. Moreover, the devices considered in the DUE algorithm should keep up with technological advances, allowing consumers to personalize their disaggregation experience even further.
- To validate disaggregation results in the absence of ground truth and explain variation in the energy behavior, the peer-comparison should be involved extensively in the analysis. Examining matching households, such as those located in similar climates, with similar numbers of inhabitants, floor surfaces, and family compositions can provide greater insights on potential energy-saving opportunities.

The following chapter focuses on transforming the obtained insights about energy consumption patterns into actions. Particularly, it introduces building occupancy forecasting algorithms that help to identify promising energy-saving opportunities that will not compromise occupants' comfort. Thus, predicting occupants' presence and absence represents a valuable input for energy management strategies aiming to reduce energy usage.

5 Occupancy Forecasting

Successful implementation of demand-side management strategies improves the grid's reliability and efficiency through timely execution of short-term energy demand adjustments on consumers' side. However, to achieve maximum impact and increase consumers' engagement, such demand-shaping actions should be accomplished without compromising consumers' comfort and convenience. Therefore, various machine learning-based occupancy forecasting algorithms are deployed to identify best-suited windows of opportunities and deliver insights on one's presence and absence patterns.

The main **highlights** and **contributions** of this chapter are:

- The development of two highly-accurate context-based approaches to forecast binary occupancy using advanced machine learning techniques. The supervised method predicts occupancy from labeled electrical energy consumption data, while the unsupervised approach employs ambient environment measurements.
- Various techniques, such as extensive feature engineering, feature selection, and improved majority voting, are introduced to enhance the performance of occupancy forecasting algorithms.
- The algorithms' performance is evaluated on a real-world dataset collected in Portugal. The supervised methodology shows the highest forecasting accuracy of 98.3% while the unsupervised approach achieves 97.6%. The importance of choosing the appropriate sensor combination and accounting for available computing resources is emphasized.

Related **publications**:

[2] Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Rule-based scheduling of air conditioning using occupancy forecasting*. Energy and AI, vol. 2, page 100022, 2020. DOI: 10.1016/j.egyai.2020.100022

5.1 Background and motivation

As discussed in Chapter 4, the shift in utilities' key role from energy provision to energy management due to increasing penetration of RES and augmenting self-consumption has facilitated the adoption of active and flexible DSM mechanisms in daily energy operations. While DSM technologies vary greatly, from load shifting and flexible load shaping to energy efficiency, there is no doubt that successful DSM activities should equally benefit energy consumers and energy suppliers in both short and long-term [Gellings 2017]. The smart building ICT framework provides diverse capabilities to understand and manage consumers' energy demand. However, when it comes to the consequent implementation of DSM mechanisms, the responsibility for success is shared between providers and consumers.

The short-term energy demand adjustments, such as time-shifting, regulation, or usage-reduction of energy-consuming appliances, are behavioral and temporary in nature. However, regardless of the chosen control type, whether automated or effectuated by consumer's action, the timeliness matters. Thus, accurate knowledge of consumers' presence in the smart building becomes interesting leverage in the successful implementation of DSM programs. The purpose of using occupancy data to take energy management decisions is threefold. First, the occupant physically present in the building is more capable of performing the necessary energy-adjusting actions in response to the DSM strategy. Second, the energy management actions aligned with the schedule of building's users or inhabitants are more likely to contribute to their well-being and comfort while increasing the potential acceptance rate of managerial actions. Third, ubiquitous digitalization with frequent notifications and actions required from the consumer's side can be overwhelming, thus demotivating the participation in DSM-related activities. Therefore, accurate knowledge of occupants' presence can ensure that requests are sent at appropriate times to deliver the maximum impact. Despite the clear benefits of occupancy-centered energy management strategies, the direct tracking of occupancy is not advised due to privacy preservation and security concerns. Thus, efficient indirect occupancy detection and forecasting methods based on advanced ML techniques have to be deployed.

The remainder of this chapter is organized as follows. Section 5.2 presents state of the art in the occupancy forecasting field, subdividing the various approaches into schedule-based and context-based. Section 5.3 introduces the proposed supervised and unsupervised occupancy forecasting methodologies. Section 5.4 describes the case study setup along with adopted performance metrics and baselines. Section 5.5 presents the results obtained by both occupancy forecasting approaches on the case study dataset. Finally, Section 5.6 concludes the chapter and proposes a three-step procedure to integrate the occupancy forecasting digital service into broader smart building ICT framework.

5.2 State of the art

Occupancy prediction algorithms can be separated into two main categories: schedule-based and context-aware. The first one represents a group of data mining methods, where the goal is to extract characteristic patterns from historical occupancy data. In the review paper [Kleiminger 2014], several algorithms were evaluated and compared to the baseline predictor. It was found that the best prediction accuracy achieved is relatively close to the estimated theoretical limit of 90% for predicting occupancy through the typical schedule derivation. Another work from [Liang 2016] demonstrated a case study in office buildings using decision trees and k-means clustering with different distance measures. The inclusion of features related to season, daylight saving time, and weekday allowed them to reach 80% of accuracy.

There are several advantages of applying a schedule-based approach. First, historical occupancy's sole usage allows minimizing the error related to the inclusion of multiple information flows and faulty sensors. Second, the possibility to work with different time resolutions gives value even to recordings with low frequency, such as one hour. However, the disadvantages associated with this methodology justify the existing theoretical limit for prediction accuracy. First, the volume of considered readings should be sufficient to obtain a reliable pattern, making data collection and labeling a challenge. Second, computational costs are very high, which is directly dependent on the first limitation. Last, constructed occupancy models are not generalizable as they represent the patterns of particular environments and occupants. As a consequence, schedule-based models cannot be applied to a different building type and can even become outdated due to the building's dynamics. As a result, context-based approaches emerged.

Context-based occupancy prediction relates to the methods that entirely depend on sensing the conditions of indoor climate or presence approximation through electricity or water usage. Depending on which sensors are used, direct and indirect approaches can be determined. Direct sensing refers to deploying devices, the purpose of which is to detect presence, such as video camera, Radio-Frequency Identification Tag (RFID), motion or sound detectors, and Passive Infrared Sensor (PIR). Most popular indirect sensors include measuring temperature, CO_2 concentration, Relative Humidity (RH), light, electricity consumption, and volatile organic compounds. Table A.3 provides an extensive characterization of various sensors used for occupancy detection, count, and tracking, while the detailed list of their advantages and disadvantages is compiled in Table A.4. As primary purposes and working principles of sensors vary, choosing the device appropriate for a specific occupancy-related task becomes a challenge. Besides, other important factors, such as installation cost, maintenance frequency, ease of data collection and processing, operational conditions, and space typology, have to be considered. Therefore, combining multiple sensors with complementary capabilities into sensor networks constitutes a practical approach in the field. Depending on which sensing devices or their combinations are used, the indirect context-based approaches for occupancy forecasting can be divided into three groups.

The first group of context-based studies includes the methods to forecast occupancy solely from electricity consumption measurements. The author in [Kleiminger 2015a] has achieved 94% accuracy by deploying supervised ML algorithms, such as Hidden Markov Model (HMM), Support Vector Machine (SVM), and K Nearest Neighbours (KNN), to extract features from raw smart meter data. Later, [Becker 2018] extended this work by proposing unsupervised algorithms that rely on 30-minute electricity consumption. The reported accuracy varied between 74% and 78%, depending on the public dataset used. Researchers in [Akbar 2015] obtained 94% accuracy using the reactive power, phase angle, and voltage and current measurements as features, thus limiting the model's utilization due to the difficulty of collecting such data at low frequencies. A simple statistical approach was introduced in [Chen 2013] to provide binary occupancy estimations. Due to the algorithm's light computational weight, it can be deployed in real-time, although the results might be erroneous at night. Authors in [Razavi 2019] used 30-min resolution smart meter measurements to detect occupancy in a large-scale household study using the gradient boosting method that led to 98% of prediction accuracy. A novel Long Short-Term Memory (LSTM) neural network approach with 90% accuracy was demonstrated by [Feng 2020] together with a convolutional neural network to detect binary occupancy in real-time from smart meter data. As the problem was framed in a supervised manner, the real occupancy was required to train the model.

The second group of context-based works agglomerates the methods that used ambient environment measurements to predict occupancy. Researchers in [Candanedo 2017] obtained accuracies higher than 97% for detecting occupancy from light, CO_2 , RH, and temperature using linear discriminant analysis and regression trees. However, the usage of the supervised approach requires retraining such models every time the sensors change location. The authors in [Yang 2012] achieved 87% self-estimation accuracy using the radial-basis function neural network by augmenting the standard ambient environment dataset with motion and sound sensors. A similar dataset complemented by light sensors was used in [Parise 2019] to predict binary occupancy in a classroom employing the SVM. The algorithm showed 96% of accuracy on a two-week case study and good scalability. A combination of temperature and motion sensors was used in [Huchuk 2019] to test various ML algorithms, such as random forest and LSTM, in residential buildings. The inclusion of time-related features allowed them to reach 80% daily average accuracy on a three-hour prediction horizon. Researchers in [Shin 2020a] applied decay function occupancy estimation based on CO_2 and motion data with 95% accuracy during a two-week experiment. A combination of occupancy detection using decision trees and prediction employing HMM was demonstrated in [Ryu 2016] on the broad set of measurements, including both indoor environmental data and energy consumption from lighting and devices. Obtained prediction accuracy varied between 85% and 93%, while model limitations were noticed in low-occupancy cases. Researchers in [Elkhouchi 2019] predicted the number of occupants solely from CO_2 concentration measurements using the LSTM model. As they used it to forecast the CO_2 evolution and not the occupants count directly, the accuracy reached was in the range of 70%.

The third group represents a significant cluster of occupancy forecasting studies driven by the proliferation of ICT. Using WiFi and Bluetooth connectivity data offers low cost, easy access, and no additional infrastructure investments. However, high risks of occupancy underestimation and overestimation remain. The prior is due to people who do not carry mobile devices, while the latter is due to various machines and appliances that are not associated with people connected to the network. The authors in [Qolomany 2018] used the WiFi beacon time series to predict occupant numbers through LSTM and ARIMA on a university collected dataset. Although difficulties in hyperparameter tuning and long training times were reported, the advantage of LSTM to overcome the vanishing gradient problem was noted. A collection of various supervised ML techniques was tested in [Rahaman 2019] on the Bluetooth sensor network. The authors achieved 90% accuracy in a commercial building using statistical features derived from raw signals. Another application of LSTM was demonstrated in [Pešić 2019] using the WiFi and Bluetooth data fusion. The attempt was made to predict binary occupancy in a residential building for a one-week horizon. Researchers have noted that changing the hyperparameters of LSTM did not affect the algorithm's performance. The authors in [Wang 2019c] detected occupancy count from WiFi with only 85% accuracy due to noise from random visitors and long-connected appliances such as printers. An interesting addition to WiFi data using computer activity measurements was proposed in [Howard 2019]. Linear regression, ANN, recurrent neural network, and LSTM models were used on one-hour resolution data and achieved 90% average accuracy for counting occupants in the office building. However, specific calibration procedure applied to collected data complicates the usage of the method. Adaptive lasso filtering for detecting occupancy levels from WiFi, CO_2 , temperature, and RH data was proposed in [Wang 2019a]. Although the 86% accuracy was achieved on a short validation experiment, the collection of media access control addresses might have compromised the occupants' privacy.

Although occupancy prediction through sensor networks generally demonstrates good results, some limitations still exist. First, sensor usage brings not only enhanced vision of the surrounding environment but additional uncertainties: dependence on sensor location and calibration procedure as well as contamination by intrinsic noise. Researchers in [Zimmerman 2018] have noticed specific phenomena related to the CO_2 sensor. Slow reaction to ambient changes exhibits a long decay of measured values, leading to misclassification between occupant's presence and absence. Second, sensor recordings represent a large volume of data, difficult to feed in the algorithm in its raw form due to the increasing demand for computing power. Therefore, careful feature engineering using domain knowledge is required to withdraw meaningful information. Finally, the deployment of a vast amount of sensors has an intrusive nature, which can compromise occupants' privacy and security. Therefore, an appropriate level of detail is needed when it comes to occupancy prediction. The complete taxonomy of solving the occupancy problem using the concept of input and output data flows is presented in Figure 5.1.

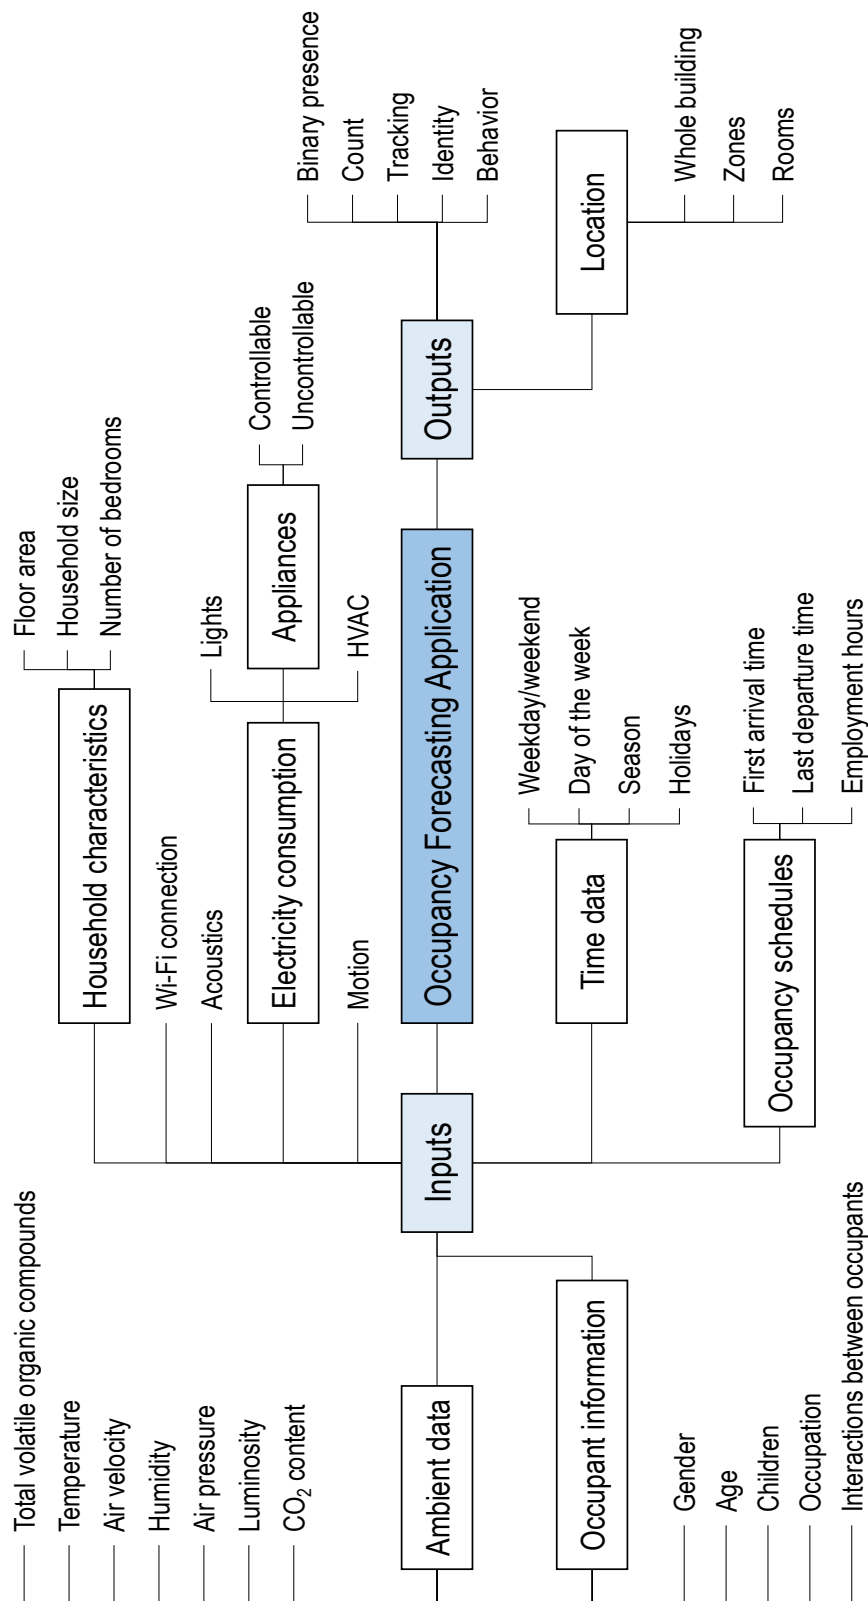


Figure 5.1 – Taxonomy of occupancy forecasting based on input and output flows

5.3 Methodology

In this section, two context-based approaches to occupancy forecasting are presented. Despite both methods being based on advanced ML techniques, the indirect sensing information utilized to infer occupancy differs. The supervised approach described in Section 5.3.1 uses labeled electrical energy consumption data, while the unsupervised one presented in Section 5.3.2 employs ambient environment measurements. The algorithms' purpose is to predict binary occupancy, while the choice in favor of indirect sensing contributes to preserving occupants' privacy.

5.3.1 Supervised method

The supervised approach aims to predict binary occupancy in a non-intrusive manner using low-frequency electrical consumption data. Thanks to proliferating deployment of smart meter technology, the availability of such data becomes widespread, thus rendering the suggested methodology suitable for both residential and tertiary sectors. The challenge of occupancy forecasting is framed as a classification problem, thus requiring the algorithm's input data to be labeled. The complete development pipeline of the supervised occupancy forecasting algorithm is depicted in Figure 5.2.

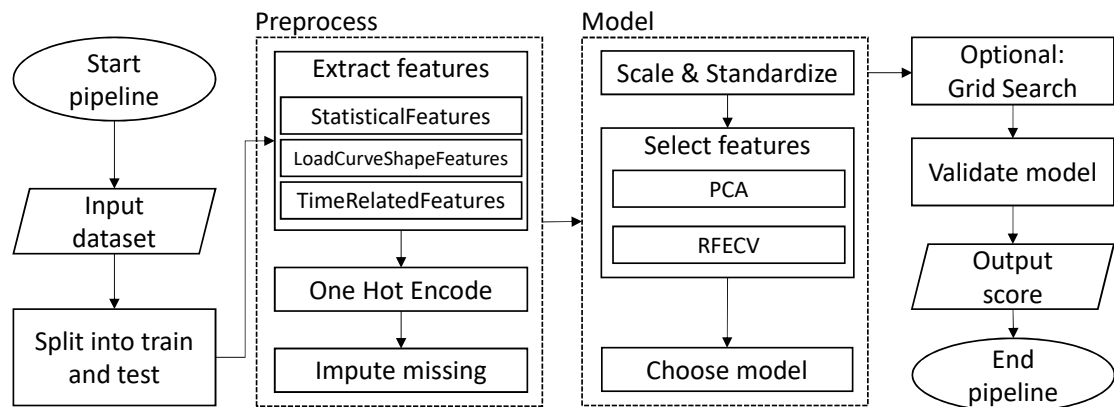


Figure 5.2 – Supervised pipeline for occupancy forecasting from electrical consumption data

The main process steps making up the pipeline include data loading, feature extraction, pre-processing, modeling, validation, and score estimation. As the proposed method is supervised, following the standard sequence of splitting the input data into training and testing sets is mandatory. Therefore, the pipeline usage allows executing the same procedure at all stages of the algorithm's development, eliminating the need to select the steps manually. The input dataset consists of smart meter power measurements and binary occupancy records, both obtained with 15-minute resolution. The latter can be either collected automatically through BMS or annotated manually. Afterward, the data are split into train and test sets, where the latter dataset is held out to provide an unbiased evaluation of the final model.

To better represent the underlying problem to predictive models, the raw electrical consumption data are transformed into a set of features describing the data. Such extensive feature engineering procedure applied to raw data in preprocessing leads to increased accuracy of occupancy prediction compared to previous works [Razavi 2019, Kleiminger 2015b, Vafeiadis 2017]. Therefore, over 60 manually designed features are created, segmented into three distinctive groups: statistical features, load curve shape features, and time-related features. The first group is based on mathematical functions, such as min, max, mean, standard deviation, median, variance, sum, and variations of their ratios. Therefore, it provides a statistical description of the load curve and gives a relative way to compare the magnitude of consumption across days. The second group consists of parameters that describe the shape of the load curve. Examples include the definition of peaks and valleys, shape similarity between same days of the week, change to the relative level of night consumption, an area under the curve, and others. This group allows capturing consumption patterns variability throughout the days, which are hypothesized to indicate changing occupancy behaviors. The third group identifies whether the measurement was taken during the weekday or a weekend, at which time of the day, month, and season it was recorded, and whether that day was a public holiday. The latter is determined based on the geographical origins of data. The underlying idea of defining this group is that consumption reflects established occupancy routines repeated through time. The full list of created features can be found in Table A.5.

To summarize, feature engineering's role is to highlight the load curve's characteristics and thus aid the inference process. Moreover, extensive feature engineering adds flexibility to the algorithm by allowing less complex machine learning models without significant losses in performance. Although statistical and time-related features are seen in previous research in different variations [Razavi 2019, Kleiminger 2015b], the definition of the load curve shape features, to the best of our knowledge, appears first in the current work. Besides, one might argue that manual feature engineering is a tedious process, and it does not necessarily produce the optimal and complete set of features. Researchers in [Razavi 2019] have thus developed a dynamic feature engineering procedure using genetic programming. However, their results did not demonstrate any increment in forecasting accuracy related to the automation deployed. Therefore, there is no evidence currently in favor of manual or automatic feature extraction.

Once the feature engineering step is complete, the successive one hot encoding step transforms categorical features into dummy binary variables. This step represents data in a uniform numerical way, which is required by most ML algorithms. The following imputation of missing values fills in boolean and float features based on most frequent and median strategies. To avoid a stronger influence of features with larger values, scaling and standardization of data are performed, converting features to zero mean and variance one. It is important to note that depending on the dataset and model being explored, not all features have equal influence on the final performance. Therefore, a feature selection step based on Principal Component Analysis (PCA) and/or Recursive Feature Elimination with Cross-Validation (RFECV) techniques is implemented. The prior technique aims to reduce dimensionality by transforming the original feature space. Thus, it creates a new set of independent features that explains 95%

of the input data variance. The disadvantage of this method is that the physical meaning of the feature space variables gets lost during the transformation. The latter technique instead keeps the original features and recursively eliminates them based on the model performance. Once the model is built and evaluated, RFECV ranks the importance of each variable in the final accuracy score. Best performing features pass to the next round of selections, while the worst performing subset is eliminated. The 3-fold cross-validation is used to determine the optimal number of best features and to avoid overfitting. This number of folds is adopted to achieve an acceptable running time of the RFECV algorithm. To choose the final feature selection technique, a set of tests where PCA and RFECV were applied both separately and in combination, was performed.

After feature selection, one has to choose the ML model. The applied preprocess and model steps, shown in Figure 5.2, are the same for all models. However, one should keep in mind that some feature selection methods, such as RFECV, require the ML model to be chosen in advance. The algorithms evaluated in this work include linear SVM, feedforward ANN with and without dropout, and ensemble methods based on decision trees, particularly Bagging and AdaBoost. The latter combine the predictions made by several base algorithms to produce the final prediction. Bagging methods do this in parallel, while AdaBoost does it sequentially. Overall, using ensemble methods allows to create estimators that are more robust and more powerful than a single estimator. The SVM model is chosen for its simplicity and extensive usage in the field of occupancy forecasting in buildings. Moreover, SVM models are known for their ability to deal with high-dimensional feature spaces and non-linear decision boundaries depending on the kernel selected. The ANN models are chosen for their flexible architecture adaptive to problems and capability to recognize non-linear relationships within data. However, difficulties in interpreting and tuning the hyperparameters are among the model's disadvantages [Dai 2020]. Also, the dropout regularization is used to reduce overfitting.

The input dataset is divided into training and testing subsets with an 80:20 ratio to evaluate the models' performance. The training subset is used to learn the inferences of occupancy from labeled electrical consumption data. To avoid overfitting, 5-fold cross-validation is used. Here the number of folds is increased to achieve less bias towards overestimating the true expected error. Instead, the testing subset is not involved in the learning process and therefore acts as unseen data to test how generalizable the model is. Further grid search can be applied to the model for finer tuning when necessary. If the model parameters have to be refined to select the optimized model, this step performs an exhaustive search over a predetermined range of values. As such procedure is computationally heavy, a trade-off between available computing resources, time at disposal, and the importance of implementing grid search must be assessed. However, optimizing models via the grid search is rarely needed, as significant accuracy improvements compared to the baseline can be achieved due to the feature engineering proposed.

5.3.2 Unsupervised method

This method aims at forecasting occupancy from ambient environment data in an unsupervised manner, which becomes useful when the supervised method is not applicable due to the absence of labeled data required for training. Such ambient environment data consist of indoor temperature, CO_2 concentration, RH, and luminosity measurements, which can be collected through sensors installed inside the building. One might argue that the availability of this kind of data is rarer than that of the smart meter measurements chosen in the supervised method. However, as the global home automation market is growing, driven by the proliferation of the IoT technology and changing lifestyles [Research 2018], ambient environment measurements will become more common. Moreover, the proposed unsupervised method can be seen as complementary to inferring occupancy from electrical consumption data. The ambient environment measurements provide an additional level of information about the building, which can enhance the forecasting performance. The methodology for predicting occupancy from ambient environment data is described in Figure 5.3.

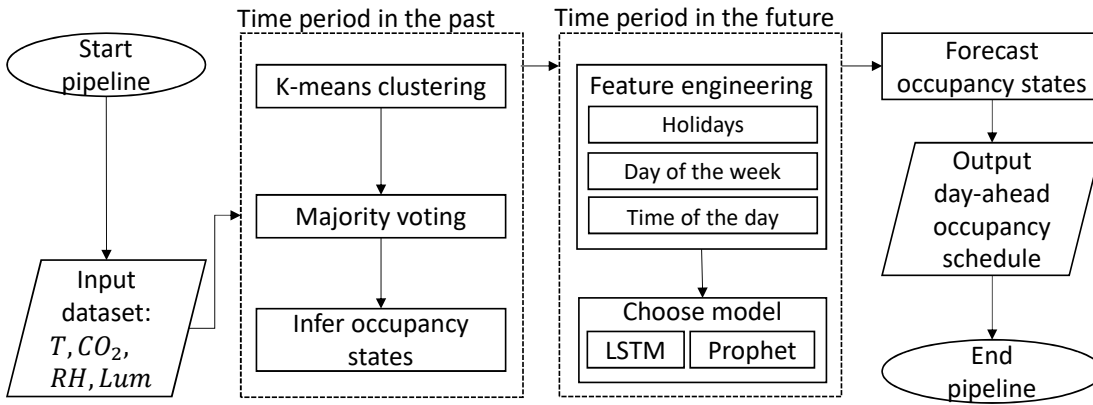


Figure 5.3 – Unsupervised pipeline for occupancy forecasting from ambient environment data

The algorithm consists of two parts that are executed in two successive periods: the time period in the past and the time period in the future, for which we are interested in forecasting occupancy. Although the future time period of interest is equivalent to one day in the current work, one can choose the duration of the prediction time window arbitrarily. Similarly, the length of the time period in the past can vary. However, it is crucial to select the dataset that will be representative of occupancy patterns and corresponding seasonalities.

First, occupancy states of the chosen time period in the past are inferred from ambient environment measurements using unsupervised K-means clustering ML technique. As the aim is to infer binary occupancy, each sensor's readings are consecutively clustered into two groups, where each observation is attributed to the cluster with the nearest mean. Figure 5.4 shows an example of such a K-means clustering procedure applied to indoor temperature data. The orange and blue colors represent two resulting clusters with the mean values indicated.

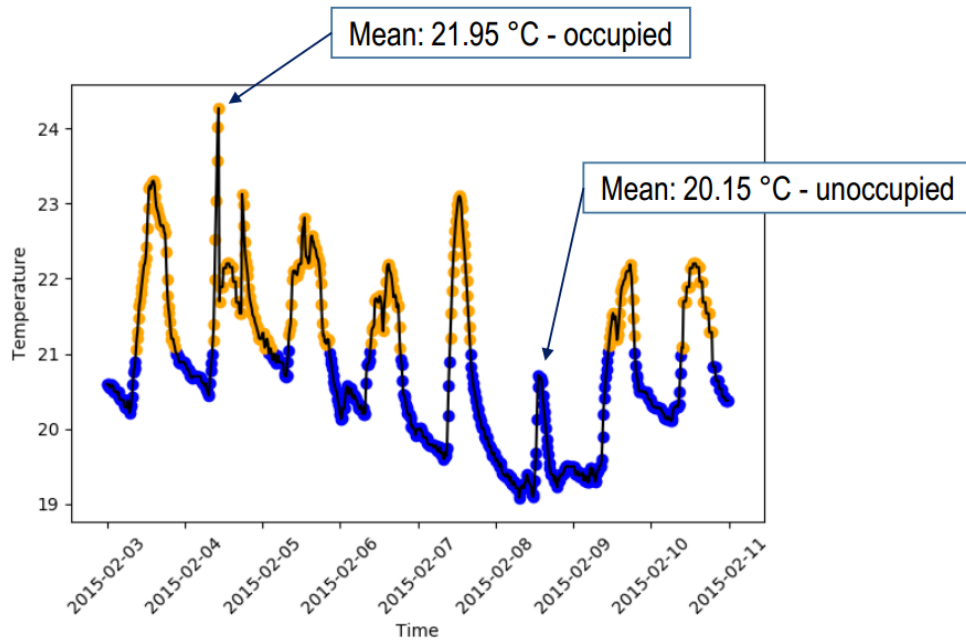


Figure 5.4 – Example of unsupervised K-means clustering technique applied to indoor temperature measurements

Once the inference of each sensor's presence and absence states is complete, the obtained information needs to be merged to provide the final decision of the occupancy state at the current time point. Although various sensor fusion techniques are suitable for this task, such as fuzzy logic, Kalman filter, and Dempster-Shafer theory, the majority voting is chosen for its simplicity and ease of interpretation. In this work, the methodology proposed in [Habib 2017] is extended by adding the trust layer that resolves situations when the decision about the room's occupancy among an even number of sensors is a draw. Although all sensors have equal decision power in the initial round, the importance of each sensor's contribution to the final choice of occupancy state is determined by its expertise. Figure 5.5 describes the second-round majority voting procedure, where particular trust weights are attributed to each of the sensors based on the number of times when the sensor's occupancy decision matched the final occupancy decision in the initial round. It is important to note that no occupancy ground truth is required as the process is fully unsupervised. The attribution of trust weights to each of the sensors supplying the input information helps to level known problems in the field such as slow decay of CO_2 concentration [Zimmerman 2018]. The impact of trust weights inclusion on the occupancy forecast accuracy will be discussed in Section 5.5.2.




Figure 5.5 – Description of the majority voting procedure

Once occupancy states in the past are inferred, powerful time series forecasting models are deployed to predict occupancy for the day ahead. In this thesis, two algorithms, namely LSTM neural network and Prophet [Taylor 2017], are applied to the problem of occupancy prediction from ambient environment data. LSTM is a particular type of recurrent neural network capable of learning long-term dependencies and thus, contrary to traditional ANN, using previously seen information to make predictions. In the current work, a specific LSTM structure is designed with two LSTM layers and 20% dropout in between to avoid overfitting. Additional optimization of batch and epoch sizes allows adjusting computational performance with respect to utilized time and resulting accuracy. Prophet is another model used for day-ahead forecasting of the occupancy patterns inferred in the past. The time series forecasting procedure developed by Facebook is based on an additive model where non-linear trends are fit with yearly, weekly, and daily seasonalities. In the model proposed, additional regressors such as weekends, national holidays, and time of the day are included to improve forecasting accuracy. A further choice between Prophet and LSTM models should be made based on the trade-off between achieved performance and required computational effort.

5.4 Case study

5.4.1 Simulation setup

The dataset, used to validate the two occupancy forecasting methods proposed, was collected with a 15-minute resolution from a university building in Porto, Portugal, in the framework of the FEEDBACK project. The collection process lasted 12 weeks that spanned from October 2018 to March 2019. To evaluate the algorithms' performance, two subsets of data, train dataset and test dataset, were created covering weeks 1 to 8 and 9 to 12, respectively. The ambient environment measurements were recorded using a multi-sensor solution developed in the course of the FEEDBACK project. The ground truth occupancy information was obtained in a non-intrusive manner, using the clock-point cards swipe count at the building's entrance doors. Figure 5.6 depicts the distribution of collected data according to the occupancy's binary classification. Although a slight skew towards presence exists in both train and test datasets, the amount of samples in each class has the same order of magnitude. Therefore, the dataset can be considered balanced.




Figure 5.6 – The dataset distribution according to 'presence' and 'absence' classes

5.4.2 Performance metrics

Different validation techniques are used for supervised and unsupervised methods to evaluate the performance of occupancy forecasting algorithms. The k-fold cross-validation followed by testing on the unseen data gives an understanding of the supervised model's capability to generalize well to an independent dataset. In the current validation procedure, the dataset is partitioned into $k = 5$ subsets and the following formulation of the accuracy score for binary classification is used:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (5.1)$$

Where True Positive (TP) signifies occupants' presence predicted as presence, True Negative (TN) is absence predicted as absence, False Positive (FP) refers to absence predicted as presence, and False Negative (FN) is the presence predicted as absence. The cross-validation score of the model is the average of the accuracy scores received at each fold. There are two reasons why accuracy was chosen to represent the occupancy forecasting performance. First, accuracy is the most common metric in the field [Dai 2020, Jung 2019]. Therefore, its usage advances the field by allowing comparison between different studies. Second, the dataset used for validating the algorithms is balanced. Thus accuracy provides a meaningful way to evaluate the performance. In addition to the accuracy score, in Section 5.5.1 the confusion matrices are reported for selected supervised models to give more details about their performance.

For unsupervised algorithms, the walk-forward validation procedure is used with accuracy scoring according to Equation 5.1. Once the occupancy is inferred on a historical dataset, the corresponding prediction is made for the day ahead. This one day constitutes the length of the walk in the validation, therefore at the next iteration, the day ahead prediction completed previously will be appended to the end of the historical dataset, and the prediction process will be repeated for another day. The overall score is the average of the accuracy scores obtained at each walk over the entire test dataset. Both validation techniques represent repetitive experiments that are completed over distinct subsets of data. As each subset population's characteristics might change, the models might exhibit different performances at each fold or each walk for cross-validation and walk-forward validation, respectively. Therefore, alongside mean accuracy, the standard deviation value for each model is provided to indicate the performances' variability across experiments.

5.4.3 Baselines

To assess the performance and models' quality, the results are compared to established baselines that differ for supervised and unsupervised occupancy forecasting models. The reason for having different baselines lies in principal differences between validation techniques used.

For supervised occupancy prediction, two specific forecasts are established as baselines: all-ones forecast and power variation forecast. The prior means that presence, which is equal to 1 in binary logic, is broadcast for all the time steps of the day ahead. The latter, instead, predicts presence in the building when the power exceeds 1.24 times the minimum power of the day. Otherwise, the absence is predicted. The 1.24 proportional constant is derived empirically for the building used in the case study as the number that realistically differentiates presence and absence according to how the load changes throughout the day. The baseline that shows the higher accuracy of prediction for the following day is used as the comparison value.

For unsupervised occupancy forecasting, the baseline is established differently. Since the walk-forward validation technique is chosen, the baseline should represent the rolling forecasts' repetitive nature, thus making a naive all ones baseline inapplicable. Therefore, the occupancy forecasted for the day ahead is calculated as the mean occupancy of the preceding week for the sake of capturing the weekday and weekend differences in occupancy. The overall accuracy of the baseline is then derived as the average of accuracy scores at each walk.

5.5 Results

5.5.1 Supervised method

This section presents occupancy forecasting results obtained using a linear SVM, Bagging, AdaBoost, and feedforward ANN with and without dropout models. The labeled electrical consumption dataset utilized in the case study is visualized in Figure 5.7 using PCA.




Figure 5.7 – A 3D representation of labeled electrical consumption data obtained using PCA

The supervised models' results are presented according to the validation procedures discussed previously. Thus, the mean prediction accuracy and its standard deviation are reported for the cross-validation phase, while for tests on the unseen data, a single prediction accuracy value is associated with each model. Table 5.1 introduces a comparison between different feature selection methods applied to supervised occupancy forecasting algorithms. The feedforward ANN is not included in this comparison as prior feature selection is not essential for this type of model due to its innate ability to develop complex relations among features. One can notice that all models benefit from the feature selection step as their prediction accuracy increases compared to when all features are fed at the input. In particular, RFECV outperforms both PCA and PCA+RFECV combination. While RFECV feature selection method is optimized based on the algorithm's final accuracy, the PCA is target agnostic and serves primarily for choosing the axes among which the input dataset can be explained best. Therefore, an inferior performance of PCA can be explained by the fact that it was not explicitly designed for feature selection.

Table 5.1 – Evaluation of feature selection methods for supervised occupancy forecasting models. On cross-validation, the mean accuracy and (standard deviation) across folds are reported, while on test, only the accuracy.

Model	None		PCA	
	cross-validation	test	cross-validation	test
Linear SVM	95.7% (1.2%)	94.8%	95.8% (1.3%)	95.1%
Bagging	94.8% (0.9%)	94.3%	95.1% (3.9%)	94.9%
AdaBoost	95.5% (0.8%)	95.2%	95.7% (1.9%)	95.2%
Model	RFECV		PCA+RFECV	
	cross-validation	test	cross-validation	test
Linear SVM	97.3% (0.8%)	97.6%	96.9% (1.0%)	96.2%
Bagging	96.2% (0.9%)	96.4%	95.2% (0.9%)	95.8%
AdaBoost	96.5% (0.8%)	96.2%	95.6% (1.4%)	95.1%

Although the supervised occupancy forecasting results onward include RFECV as the feature selection step, one should always be aware of what such choice implies in practice. Notably, the ranking among features and their optimal number will vary at different runs of the supervised pipeline due to the inherent stochasticity of the cross-validation procedure and ensemble models in particular. Moreover, the results will differ if the training dataset changes in any way or another supervised model is chosen for occupancy forecasting. Thus, similarly to scaling and standardization, one should reapply the feature selection procedure at every fit of the model to the training data. Nevertheless, observing the features selected across linear SVM, Bagging, and AdaBoost models highlights that certain features are chosen more consistently. Particularly, the weekday or weekend, time of the day, is peak, first order difference, and mean hourly features appear more frequently, which can be indicative of their high contribution to the forecasting accuracy.

Table 5.2 summarizes the simulation results for supervised occupancy forecasting. Additionally, the impact of applied feature engineering on prediction accuracy is demonstrated. The baselines utilized for evaluating the models' performance exhibit different behavior. The all-ones forecast, baseline 1, is slightly better than a random guess, indicating that the data is slightly skewed towards presence as shown in Section 5.4.1. The power variation forecast, baseline 2, demonstrates higher accuracy due to its more sophisticated definition. Therefore, the latter is used to evaluate the performance of other models. The overall results confirm a positive influence of the applied extensive feature engineering as it improves the prediction accuracy by 15% on average for all tested models. Moreover, it is shown that even basic usage of ML algorithms gives better results than empirically derived rules to infer occupancy from electrical consumption data. Another observation is the relatively low values of accuracy demonstrated by ensemble models, Bagging and AdaBoost, compared to the linear SVM and feedforward ANN.

Table 5.2 – The impact of feature engineering on the results of supervised occupancy forecasting. On cross-validation, the mean accuracy and (standard deviation) across folds are reported, while on test, only the accuracy.

Model	cross-validation	test
Baselines		
Baseline 1	53.7% (2.1%)	54.2%
Baseline 2	73.9% (1.9%)	72.5%
With feature engineering		
Linear SVM	97.3% (0.8%)	97.6%
Bagging	96.2% (0.9%)	96.4%
AdaBoost	96.5% (0.8%)	96.2%
FF ANN	98.6% (0.2%)	98.1%
FF ANN + drop	98.7% (0.2%)	98.3%
Without feature engineering		
Linear SVM	83.3% (3.1%)	80.1%
Bagging	80.7% (3.0%)	78.2%
AdaBoost	81.5% (2.6%)	79.1%
FF ANN	83.7% (0.7%)	82.3%
FF ANN + drop	83.8% (0.6%)	82.6%

The underlying concept of ensemble models is the aggregation of weak classifiers, typically decision trees. Therefore, the final prediction is a combination of several models trained in parallel. However, decision trees' utilization does not set the best boundaries between two different classes due to no natural conditions of occupancy derivation from load curves that would govern decision tree splits. Despite the performance being lower than expected, the AdaBoost model shows better stability when executed on unseen data due to the application of weighted averages for choosing the most contributing trees and tracking prediction errors. The linear SVM model with $C = 1$ demonstrates comparable performance to the feedforward ANN due to the advantages of extensive feature engineering and subsequent feature selection.

However, one can notice that the performance of linear SVM is less consistent as the standard deviation across folds is higher than the one for feedforward ANN. Nonetheless, the overall value of standard deviation does not surpass 5%, thus making such a performance acceptable.

Figure 5.8 depicts confusion matrices obtained on test dataset for linear SVM, AdaBoost, and feedforward ANN with dropout as they represent three families of supervised algorithms used in this study. Since occupancy forecasting aims to foster better energy management decisions and realize potential energy savings, the low values of FP are required to avoid energy waste. One can observe that for all models the number of FP is lower than FN. The FN count is directly related to the occupants' thermal discomfort, as the occupied space is predicted unoccupied. Although minimizing thermal discomfort is equally important to achieving energy savings, in reality, reaching comfortable conditions for all occupants is challenging as they vary depending on one's clothing and metabolism. Therefore, decreasing FN comes at the second priority compared to realizing potential energy savings.

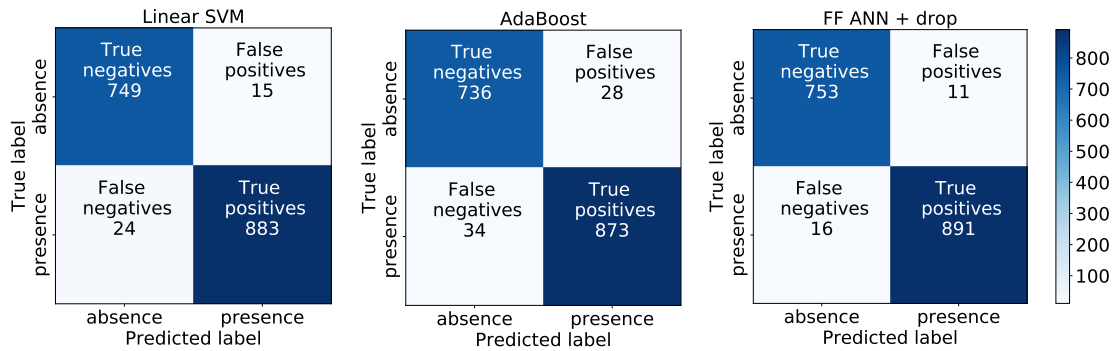


Figure 5.8 – Confusion matrices on test data for linear SVM, AdaBoost and feedforward ANN with dropout models

Precision constitutes another important indicator that can be computed from confusion matrices depicted in Figure 5.8 and is determined as follows:

$$Precision = \frac{TP}{TP + FP} \quad (5.2)$$

In the context of building energy management, the precision metric can be related to the fraction of energy consumed by the building when it was actually occupied and thus has to be maximized. If the occupancy forecasting is applied for building management with objectives different to energy savings, another interpretation of the precision metric detailed in Equation 5.2 has to be found. The feedforward ANN with dropout demonstrated superior precision value than other models: 98.8% against 98.3% for SVM and 96.9% for AdaBoost. Thus, one can hypothesize that using feedforward ANN for occupancy forecasting would result in less energy waste.

Figure 5.9 depicts the results of an optional grid search step applied to the linear SVM model. The optimization procedure aimed at tuning model's hyperparameters has confirmed that chosen value of the regularization parameter $C=1$ leads to the highest prediction accuracy. Similarly, the grid search can be applied to other models used in this study. Particularly, the number of decision trees can be tuned for both Bagging and AdaBoost models. At the same time, the size of the bootstrap sample and the learning rate can be additionally selected for the prior and the latter, respectively. For the feedforward ANN, one can choose to tune the number of epochs, learning rate, batch size, and the number of neurons in the hidden layer. As for all these models multiple hyperparameters have to be considered simultaneously, the grid search optimization procedure goes beyond one-dimensional space and thus requires substantial computational efforts to be executed.

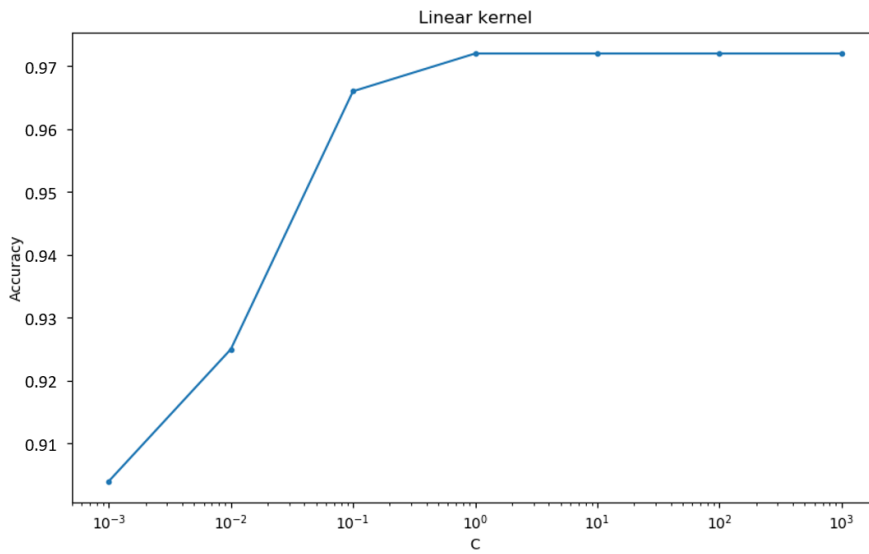


Figure 5.9 – Grid search optimization of the linear SVM's prediction accuracy as a function of regularization parameter C

To summarize the performance of supervised occupancy forecasting algorithms, one can choose simple ML models to achieve satisfactory prediction accuracies if all necessary steps are taken to aid the algorithm see the underlying structures in the input data. The accuracies achieved are higher than those reported in the literature for occupancy forecasting from electricity consumption in multi-occupancy spaces [Jung 2019].

5.5.2 Unsupervised method

Table 5.3 summarizes the results of unsupervised occupancy forecasting from ambient environment data, namely indoor temperature, RH, luminosity, and CO_2 concentration. The performance of chosen unsupervised models, LSTM and Prophet, is presented in the form of the mean and standard deviation of prediction accuracy during the walk-forward validation and compared to the baseline model. Besides, the impact of the trust layer inclusion on the accuracy prediction score is demonstrated.

Table 5.3 – Accuracy and (standard deviation) of unsupervised occupancy forecasting models

Model	no trust layer	trust layer
Baseline	45.8% (0.08%)	45.9% (0.08%)
Prophet	56.8% (0.09%)	57.1% (0.07%)
LSTM	95.7% (0.03%)	97.6% (0.01%)

One can notice that including the trust layer into the majority voting procedure is more beneficial for LSTM than for other models as it improves the prediction accuracy score of the LSTM model by 2%. Simultaneously, almost no effect is seen on the baseline, and the impact on the Prophet model is moderate. The trust layer's main goal is to verify the reliability of sensor measurements and improve the quality of the consensus reached by the majority voting. However, if the algorithm fails to infer occupancy patterns correctly, slight improvements in input data quality will not influence the overall performance significantly. Although Prophet shows some predictive power by surpassing the baseline, LSTM exhibits the highest accuracy among unsupervised models. The underlying principles of the models can explain these performance differences. As Prophet is an additive model, it fits non-linear trends and works best with strong seasonal effects. In the occupancy forecasting case, the weekly patterns, such as increased occupancy throughout weekdays, and daily patterns that correspond to two activity peaks - before and after lunchtime, are experienced. Thus, if Prophet fails to identify and distinguish those two types of seasonalities, the prediction becomes inaccurate. For both models, additional regressors such as holidays are used to explain sudden changes in occupancy patterns.

Although LSTM model demonstrates superior performance to Prophet, the final choice of the unsupervised forecasting algorithm should consider the execution time along with prediction accuracy. Moreover, as LSTM model belongs to the family of ANNs, one should analyze the hyperparameters' influence on the model's performance. Therefore, Figures 5.10 and 5.11 depict the number of epochs and batch size tuning results, respectively, alongside execution times. Obtained values describe the model's performance on a 3-week test dataset using a walk-forward validation technique. According to ANN's terminology, the number of epochs defines the number of complete passes through the training set, while the batch size refers to the number of samples processed before the model's update. Both algorithms are programmed in Python language and are executed on a personal laptop (Intel i7- 7600, 16 GB RAM).

As can be seen in Figure 5.10, LSTM's execution time depends linearly on the number of epochs and increases greatly from approximately 24 minutes to 110 minutes for 10 and 60 epochs, respectively. At the same time, the prediction accuracy increments steadily up to 30 epochs and stagnates afterward. A trade-off between prediction accuracy and execution time suggests that a marginal 0.03% accuracy gain per epoch requires a significant execution time increase of 116 seconds. Therefore, to use rationally available computer resources, the number of epochs chosen is 10.

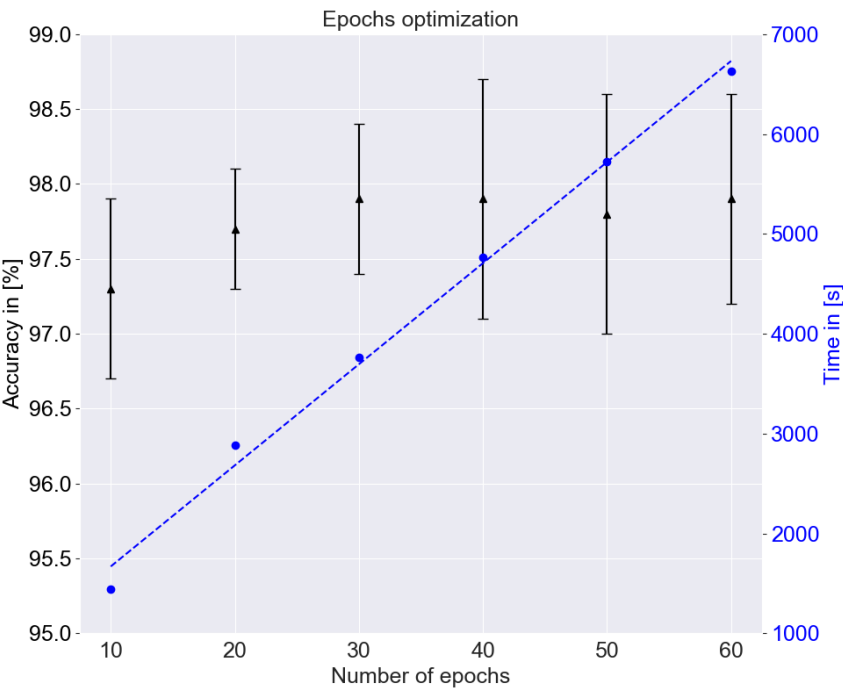


Figure 5.10 – LSTM's epoch size optimization reported along with the model's execution time. All experiments were conducted with a fixed batch size of 32, which is a default value for ANNs in Keras Python library.

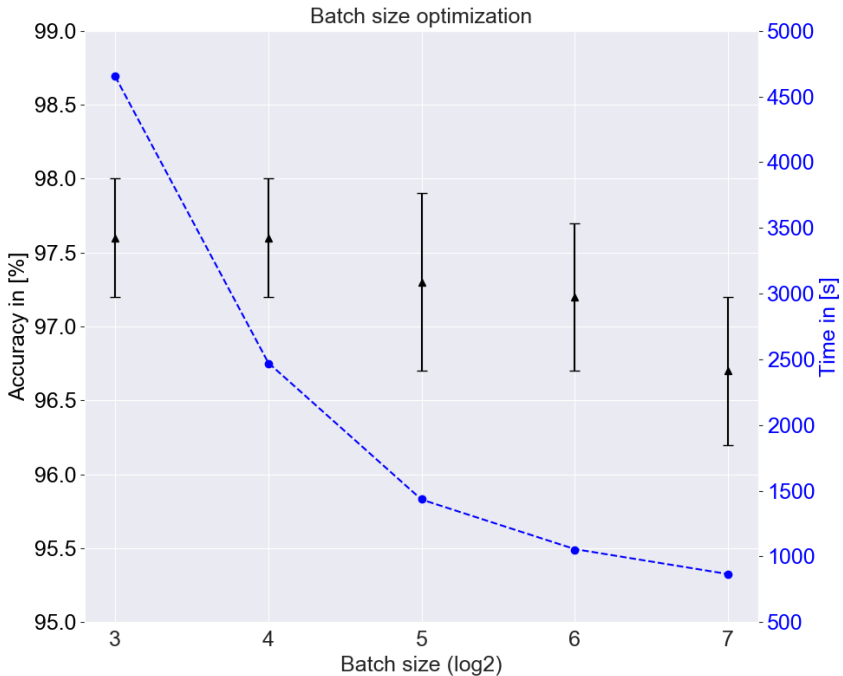


Figure 5.11 – LSTM's batch size optimization reported along with the model's execution time. All experiments are conducted with fixed number of epochs equal to 10.

The dependence of execution time on the batch size depicted in Figure 5.11 differs from the relation with the number of epochs by exhibiting the exponential nature. Particularly, both execution time and prediction accuracy reduce with the batch size increase. Although batch sizes of 8 and 16 yield the best prediction accuracy, their respective execution times are significantly larger than those of batch sizes 32 and 64. Thus, considering both execution time and prediction accuracy, the batch size of 64 is chosen as 0.1% compromise on accuracy reduces the execution time by almost 27% compared to batch size 32. Comparison of the chosen LSTM model with Prophet results in the latter executing approximately 10 times faster. However, as it was demonstrated in Table 5.3, the resulting prediction accuracy of the Prophet model is significantly inferior. The overall accuracy values reported for unsupervised algorithms are consistent with those achieved in literature for occupancy forecasting from environmental data for multi-occupancy spaces [Jung 2019].

5.6 Conclusion

To conclude, this chapter presented two context-based approaches to forecast occupancy using advanced ML techniques. The proposed supervised and unsupervised methods yielded the highest prediction accuracies of 98.3% and 97.6%, respectively, thus proving their reliability as the foundation of occupancy-centered digital energy services. To benefit from occupancy forecasting when implementing energy management strategies, one should follow the three-step procedure to successfully integrate it into the broader smart building ICT framework. First, taking into account the diversity of available sensors for occupancy detection, one should choose the appropriate sensor combination depending on building's characteristics, complexity of layout, and functional purpose. To yield maximum benefits from occupancy detection, the desired level of detail has to be considered alongside occupants' privacy. Second, the choice of sensors and availability of historical occupancy records should dictate which forecasting methodology has to be implemented. While the unsupervised method can flexibly accommodate various sensors with time-series output without any modifications in the prediction logic, the supervised technique will require additional feature engineering to incorporate anything besides smart meter measurements. Last, as demonstrated by the unsupervised case study, the choice of occupancy forecasting model depends largely on the availability of computing resources and desired level of prediction accuracy. Thus, one should carefully consider the needs of the targeted digital energy service to understand which accuracy gains and which computational speeds would make a difference in the success of occupancy-centered energy management strategies.

The following chapter focuses on deploying occupancy forecasts for intelligent building energy management that does not compromise occupants' comfort. Particularly, it introduces an occupancy-centric rule-based automation algorithm that aims to reduce the energy usage of building's heating and air conditioning systems.

6 Building Automation

The growing energy use across major sectors poses a great concern on the future development of climate change issues. Buildings represent a massive potential for energy savings, thus endowing them with an intelligent control layer can significantly improve their operational efficiency and lead to energy conservation. Being the largest energy consumer in both residential and tertiary buildings, the Heating, Ventilation, and Air Conditioning (HVAC) systems become the main target of building automation strategies that aim to increase energy efficiency.

The main **highlights** and **contributions** of this chapter are:

- The introduction of a novel occupancy-centric rule-based HVAC automation algorithm. The suggested approach employs occupancy forecasting methods introduced in Chapter 5 to generate day-ahead ON/OFF schedules for cooling and heating systems.
- The HVAC automation algorithm is evaluated on two real-world datasets collected in Portugal. The results demonstrate that potential energy savings of 15.4% and 25.7% for 6-month and 15-month periods, respectively, can be achieved.
- The seasonal analysis of achieved results shows that reducing the energy spent on cooling as opposed to heating has a higher contribution to potential energy savings.

Related **publications**:

[2] Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Rule-based scheduling of air conditioning using occupancy forecasting*. Energy and AI, vol. 2, page 100022, 2020. DOI: 10.1016/j.egyai.2020.100022

[5] Marina Dorokhova, Fernando Ribeiro, António Barbosa, João Viana, Filipe Soares, Nicolas Wyrsh. *Real-world implementation of an ICT-based platform to promote energy efficiency*. Energies, 14(9), 2416, 2021. DOI: 10.3390/en14092416

6.1 Background and motivation

Energy use across all sectors, including industry, buildings, and transportation, has grown significantly over the last decades. As buildings alone account for 38% and 40% of the total final energy consumption in the EU and the US, respectively [U.S. EIA 2020, EEA 2020a], monitoring, controlling, and conserving energy in buildings becomes paramount to efficiently respond to the challenges of the new energy era. Buildings' energy consumption is mainly influenced by several factors, such as their location, thermal properties, construction characteristics, occupants' behavior, and quality of the Heating, Ventilation, and Air Conditioning (HVAC) system. Although minimizing buildings' consumption can be done at the design stage through extensive energy simulations, the study [Delzendeh 2017] has shown that a significant gap between planned and actual energy use remains. Thus, to unlock the massive potential for energy savings, significant improvements of buildings' operations are necessary, which can be achieved by endowing buildings with an intelligent control layer.

With BMSs covering various functions within buildings such as lighting, security access, and control of electrical equipment, focusing on HVAC automation remains a priority. HVAC is the largest energy user in both tertiary and residential sectors, accounting for approximately 50% of the final energy consumption, with 75% of heating and cooling needs being satisfied by fossil fuels [Commission 2020]. Therefore, optimizing HVAC activities can bring substantial cost benefits and resolve major shortcomings, such as assuming maximum occupancy in spaces, conditioning empty spaces, and operating regardless of occupants' perspectives [Jung 2019]. As buildings' primary purpose is to serve the users and ensure their well-being, implementing occupancy-centered building automation strategies becomes important, especially with people spending 87% of their time indoors [Klepeis 2001]. Thus, the NILM technique proposed in Chapter 4 and ML-based occupancy forecasting algorithms developed in Chapter 5, can provide considerable support to understanding which energy management actions have to be taken and when, thus leading to the successful implementation of efficient HVAC control strategies that ensure occupants' thermal comfort and contribute to energy efficiency.

The remainder of this chapter is structured as follows. Section 6.2 summarizes state of the art in the field of intelligent HVAC control. Section 6.3 describes the suggested occupancy-centric rule-based HVAC control methodology and elaborates on the integration of occupancy forecasting techniques proposed in Chapter 5. Section 6.4 defines the case study, particularly focusing on the simulation setup, performance metric, and baseline. Section 6.5 presents the results achieved on two simulations and discusses the factors influencing the performance. Finally, Section 6.6 concludes the chapter and highlights the specificity of deploying the building automation digital energy service in the real world.

6.2 State of the art

Intelligent HVAC control strategies vary in their nature, complexity, and mode of operation. According to the classification provided in [Wang 2008], most control methods can be divided into two categories: local and supervisory control. The prior group of control methods corresponds to low-level management, where HVAC system's components are targeted directly to ensure safe, robust, and adequate building operations while considering dynamic characteristics of the local environment. Local control methods can be further subdivided into sequencing and process control, with prior being responsible for switching the components on and off, and the latter aiming to bring their states to desired values. The supervisory control methods, which can be either model-based or model-free, address HVAC operations from a high-level perspective utilizing global optimization techniques. Particularly, they aim to determine the optimal way to manage HVAC systems in a cost-efficient or energy-efficient manner, depending on the chosen objective function.

One of the first works in the field dates back to 1997 when researchers in [Zhou 1997] introduced a neurothermostat concept. Framing a predictive control task as an optimization problem, they defined two overarching objectives: minimizing energy costs and maximizing occupants' comfort, which became the most common goals of intelligent HVAC control. Since then, the realization methods expanded significantly, embracing the rapid progress of the ML field. Authors in [Peng 2018] proposed a demand-driven cooling strategy with learning capabilities that relies on the information retrieved from temperature and motion sensors. The algorithms, based on K-means and KNN techniques, were deployed in 11 different locations and resulted in energy savings ranging from 7% in single-person offices to 52% in multi-person rooms. Researchers in [Dong 2014] demonstrated a nonlinear Model Predictive Control (MPC) strategy that considers behavior patterns of building occupants and local weather forecasts. In February and July, they achieved 30.1% and 17.8% consumption reductions, respectively, discovering a higher savings potential for heating as opposed to cooling. Mixed-Integer Linear Programming (MILP) suggested in [Jindal 2018] has proved to be a successful technique to solve the HVAC control problem, allowing the inclusion of additional constraints, such as rooms' size. The proposed heuristic-based algorithm led to 21.2% load reduction and 16% cost savings in the university building. Researchers in [Yu 2020] implemented a multi-agent deep RL algorithm to reduce HVAC system's operational cost while maintaining occupants' thermal comfort and indoor air quality in a commercial building. Considering electricity prices, outdoor temperature, and zone occupancy, they demonstrated 75.2% and 56.5% cost reductions compared to rule-based and heuristic control. Other approaches to problem-solving in HVAC optimization field include feedback-based control, game theory, and intelligent agents, and can be generally allocated to one of the following categories: environmental condition-based, schedule-based, occupancy-based, and CO_2 -based. An extensive review of HVAC control methods conducted in [Mirakhorli 2016, Jung 2019] highlighted the importance of including additional features, such as occupancy patterns, outdoor climate, and building characteristics, for obtaining high potential energy-savings in controlled HVAC operations.

6.3 Methodology

To solve the HVAC scheduling problem, one can choose between optimization and heuristic methods. Although both of these approaches offer certain advantages and disadvantages, a rule-based heuristic method is proposed in the current methodology for the following reasons. First, the aim is to provide an improved schedule for HVAC operations in the day ahead. Therefore, the forecasts of selected variables among the input data are used, bringing significant uncertainties to the final performance. As a consequence, a solution obtained via optimization cannot claim to be optimal when exact forecasts are not available. The rule-based technique instead appears to be more robust and thus may deliver a better performance. In practice, it gives general operational guidelines which are less sensitive towards how close the actual values follow the forecast. Second, heuristic methods are less sophisticated than optimization and typically take less time to execute. This trait can be viewed as an advantage when the scheduling algorithm has to run more frequently and switch from day-ahead to intraday operations. Therefore, to minimize the impact of forecast uncertainties and achieve acceptable execution time, the following rule-based ON/OFF scheduling algorithm is proposed.

The necessary input information to produce a rule-based schedule includes forecast values of outdoor temperature, solar radiation, and occupancy. The accuracy of the latter substantially contributes to ensuring occupants' thermal comfort by providing timely heating or cooling. The methodology of HVAC ON/OFF scheduling is based on the calculation of building's thermal load $\dot{Q}_{th}(t)$ according to Equation 6.1:

$$\dot{Q}_{th}(t) = A_{th}[k_{th}(T_{int} - T_{ext}(t)) - k_{sun}\dot{i}(t) - b(t)\dot{q}_p(t)] \quad (6.1)$$

Where:

- A_{th} is the building's heated surface [m^2]
- T_{int} is the indoor comfort temperature set-point [$^{\circ}C$]
- $T_{ext}(t)$ is the outdoor temperature [$^{\circ}C$]
- k_{th} is the heat transfer coefficient representing building's thermal losses to the environment [$W/m^2/K$]
- k_{sun} is the heat transfer coefficient representing thermal gains from solar radiation [-] that considers building's shape, the fraction of window surface, and the transmittance of the glass
- $\dot{i}(t)$ is the solar global radiation [W/m^2]
- $b(t)$ is the binary occupancy profile [-], thus, when the space is unoccupied $b(t) = 0$, and no contributions to heat gain from people are taken into account
- $\dot{q}_p(t)$ is the coefficient representing internal heat gain due to people's presence that depends on the building's functionality [W/m^2]

If k_{th} and k_{sun} coefficients are absent from the building's envelope description, one can infer their values from historical data of heating system operation using the Newton-Raphson method [Girardin 2012]. As maintaining the building's indoor temperature around the chosen set-point temperature is required to ensure occupants' thermal comfort, Equation 6.1 assumes that external heating or cooling is applied as compensation to the heat losses and gains that occurred. Therefore, the procedure to obtain the values of k_{th} and k_{sun} unfolds as follows:

1. Set the heating cut-off temperature T_{off}^{heat}
2. Set the initial value of k_{th}^0 in the range from 1 to 10
3. Calculate the missing value of $k_{sun}(k_{th}^0)$ from Equation 6.1, assuming that $\dot{Q}_{th}(t) = 0$ if $T_{ext}(t) = T_{off}^{heat} \pm 1^\circ C$ and taking the mean values of solar global radiation and heat gain from people
4. Define the target function $f = 0$, where the first and second terms represent the calculated and historical annual heating demands, respectively:

$$f = \sum_{n=1}^N Q_{th}(t) - Q_{th}, \text{ where } N \text{ is the number of corresponding periods}$$

5. Calculate the new value of k_{th} using the Newton-Raphson method:

$$k_{th}^1 = k_{th}^0 - \frac{f}{f'}$$

6. Repeat until the maximum number of iterations is reached or $|f| \leq \epsilon$, where ϵ is the chosen tolerance in the range from 10^{-4} to 10^{-6}

After computing the building's thermal load and estimating the missing coefficients, the HVAC operation scheduling takes place according to the set of predefined rules. As the suggested rules are independent from each other and there is no priority among them, the whole set has to be satisfied:

- Heating ON if $\dot{Q}_{th}(t) > 0$, cooling ON if $\dot{Q}_{th}(t) < 0$
- Heating ON if $T_{ext}(t) < T_{off}^{heat}$, cooling ON if $T_{ext}(t) > T_{off}^{cool}$
- Heating ON or cooling ON if $b(t) \neq 0$, else both OFF
- No simultaneous heating and cooling operation
- Delay heating/cooling if $T_{ext}(t)$ evolution is short-term or fluctuating

The last rule prevents alternate ON and OFF switching of either heating or cooling systems. Employing this rule contributes to keeping comfortable indoor temperature if there is any ramp-up time specified for the HVAC system. Additionally, it helps to extend the lifetime of the utilized equipment and to avoid unnecessary maintenance.

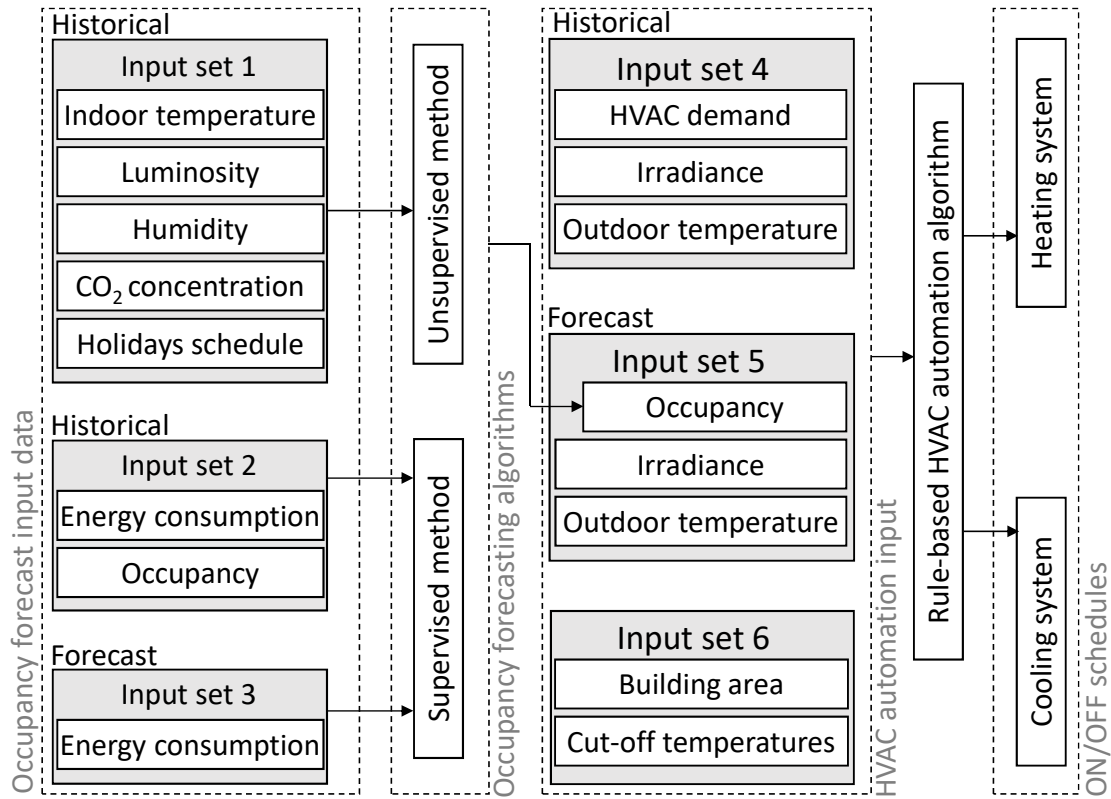


Figure 6.1 – Occupancy-based HVAC scheduling model block-scheme

Figure 6.1 represents the block-scheme of the occupancy-based HVAC scheduling model that is comprised of two main parts: occupancy forecasting algorithms and HVAC automation algorithm. As the model is sequential, the occupancy forecasting is done prior to ON/OFF schedule generation. For each of the parts, the respective input sets are identified, and corresponding data flows are indicated to describe the data utilization in the overall model. Sets 1 to 5 have either historical or forecast property to distinguish between the data collected using existing hardware and the data that emerged from predictive calculations. The input set 6, instead, contains the data constant over time and represents related building metadata. In the current work, the respective simulations are conducted at the building level. However, one has to note that subject to data availability, occupancy forecasting and HVAC automation algorithms can be applied at zone and room levels without changes in the proposed methodology.

6.4 Case study

6.4.1 Simulation setup

The dataset, used for validation of the proposed occupancy-based HVAC scheduling model was collected within the FEEDBACK project from a university building in Porto, Portugal.

All the necessary input data is divided into sets, indicated in Figure 6.1. The input sets 1 and 2 were collected with the 15-minute resolution, while input set 4 with hourly resolution. The ground truth occupancy information in the input set 2 was obtained in a non-intrusive manner, using the clock-point cards swipe count at the building's entrance doors. The input set 6 contains building's metadata. The building area A_{th} is equal to 4000 m². The cut-off temperatures for HVAC regulation are chosen according to the EN15251 comfort monitoring standard and are equal to $T_{off}^{heat}=16^{\circ}\text{C}$ and $T_{off}^{cool}=26^{\circ}\text{C}$ for heating and cooling, respectively. The HVAC automation algorithm aims to maintain the indoor comfort temperature set-point $T_{int} = 21^{\circ}\text{C}$.

In this thesis, two independent simulations are conducted to validate the proposed HVAC automation methodology. The difference between the two consists in the period used for collecting input sets 1, 2, and 4. The first simulation utilizes the same dataset as presented in Chapter 5, thus the 12 weeks spanning from October 2018 to March 2019. The second simulation encompasses a longer period from July 2019 to September 2020.

6.4.2 Performance metric

The performance of the HVAC ON/OFF scheduling algorithm is assessed based on potential energy savings E_s , calculated according to Equation 6.2:

$$E_s = \frac{\sum_{i=1}^N P_{HVAC}^i (\alpha_{conv}^i - \alpha_{impr}^i)}{\sum_{i=1}^N P_{HVAC}^i \alpha_{conv}^i} 100\% \quad (6.2)$$

Where N is the number of the time intervals in the period of interest, P_{HVAC}^i is the HVAC power consumption at time interval $i \in \{1, 2, \dots, N\}$. Conventional $\alpha_{conv}^i \in \{0, 1\}$ and improved $\alpha_{impr}^i \in \{0, 1\}$ HVAC schedule indicators represent whether the HVAC system should be turned ON or OFF at each particular time interval. Therefore, potential energy savings arise from the differences in HVAC energy consumption governed by the conventional and improved operation schedules.

6.4.3 Baseline

The baseline for HVAC automation corresponds to the conventional ON/OFF HVAC scheduling before any rule-based algorithm is implemented. According to the established schedule, in the present case study, HVAC operates uninterruptedly during the weekdays from 08h00 in the morning to 19h00 in the evening. Thus, the baseline of energy consumption is computed as the total amount of energy spent for HVAC purposes during the period of interest within daily operational time intervals.

6.5 Results

In the first simulation, the performance of the proposed HVAC scheduling algorithm was evaluated on the whole dataset that spans from October 2018 to March 2019. The unsupervised occupancy forecasting method, explained in Chapter 5, was used to predict occupancy as it does not require any training and thus can extend the evaluation period of the algorithm. A typical week in the case study consists of 5 days, Monday to Friday, according to the opening schedule of the university building.

Figure 6.2 depicts the results of a comparison between conventional and improved rule-based HVAC ON/OFF schedules for the duration of the evaluation period. As there is no prior knowledge on particular timings of heating and cooling systems in the conventional scenario, the system's total operational time frames are shown in yellow. While the improved scenario can differentiate the heating and cooling processes, in the chosen evaluation period, heating prevails as the season changes from mid-autumn to early spring. Minimization of energy spent on heating becomes more evident in October and towards the end of March due to augmented outdoor temperatures and higher irradiance values. Notably, weekly average outdoor temperatures have risen by 11% and 31% during weeks 3 and 4 of March, respectively, compared to the beginning of the month. Similarly, total solar irradiance received throughout the week has increased by 24% at the end of March. Occupancy patterns, instead, did not vary throughout the period of consideration and are representative of the typical daily occupancy regardless of the season. Such stability in the present case study is explained by distinctive office hours and the global characteristic of occupancy attributed to the whole building. Therefore, the main impact on potential energy savings comes from the rule-based nature of the algorithm.




Figure 6.2 – Simulation 1, October 2018 - March 2019: comparison between conventional and improved ON/OFF schedules of the HVAC system (segment's width = ON or OFF duration).

During the evaluation period, a potential of 15.4% energy savings is demonstrated when switching to improved scheduling. This value is consistent with median energy savings reported in literature for occupancy-reactive HVAC control [Jung 2019]. The highest potential savings are achieved during the spring season with a relative share of 42%, while winter and autumn seasons have contributed to savings with 26% and 32% shares, respectively. However, as data collection was not constant throughout the seasons, one should compare the values of relative daily savings. Thus, in autumn, daily potential energy savings correspond to 2.7% on average, while in spring and winter, the values are 2% and 0.9%, respectively. The small amount of potential energy savings in winter can be explained by low outdoor temperatures and low irradiance values. The average outdoor temperature in winter is 5 degrees less than in autumn and 1 degree less than in March. At the same time, the average solar irradiance values in winter have fallen by 30% and 44% compared to autumn and spring, respectively. Although insufficient heat influx from outdoors makes heating demand in winter inevitable, the summer HVAC consumption in Portugal is typically three times higher. Thus, further research is required to evaluate the rule-based algorithm's potential to achieve HVAC consumption reductions in summer. Moreover, it is expected that occupancy will have a higher impact on energy savings in buildings with variable occupancy patterns and increased spatial resolution.

To test this hypothesis, the second simulation encompassing a longer period from July 2019 to September 2020 was conducted. The difference between the two simulations is twofold. First, the chosen time period contains two summer seasons, which gives a unique opportunity to evaluate potential energy savings resulting from cooling. Second, the dataset is partially collected during the COVID-19 pandemic [WHO 2021], which introduces disturbances in established occupancy patterns and results in higher occupancy variability. As a consequence, the SVM-based supervised occupancy forecasting algorithm described in Chapter 5 was used to predict occupancy. The underlying reason is that utilization of consumed electrical energy provides a more robust input as opposed to ambient environment measurements, which depend largely on sensors' indoor localization and can mislead occupancy forecasting by reflecting natural temperature, RH, and luminosity cycling. Similar to the first simulation, the weekends were removed from the dataset to account for the university building's typical opening schedule. Moreover, to consider the effects of COVID-19 and identify specific periods when the university was supposedly closed, the days where less than 10% of occupants were present according to the ground truth were equally removed. Therefore, out of 335 days recorded initially, 101 days represent weekends and 66 days constitute low-occupancy days.

Figure 6.3 depicts the comparison between conventional and improved HVAC ON/OFF schedules produced for the dataset of Simulation 2. The heating and cooling periods are expectedly distributed according to seasons, although one can notice some short heating instances at the beginning of summer 2020. The reasons for that are low morning temperatures below the heating cut-off combined with poor solar irradiance. During the observation period, the implementation of the rule-based HVAC automation algorithm resulted in 25.7% of potential energy savings, which are detailed by season and HVAC functionality in Tables 6.1 and 6.2.

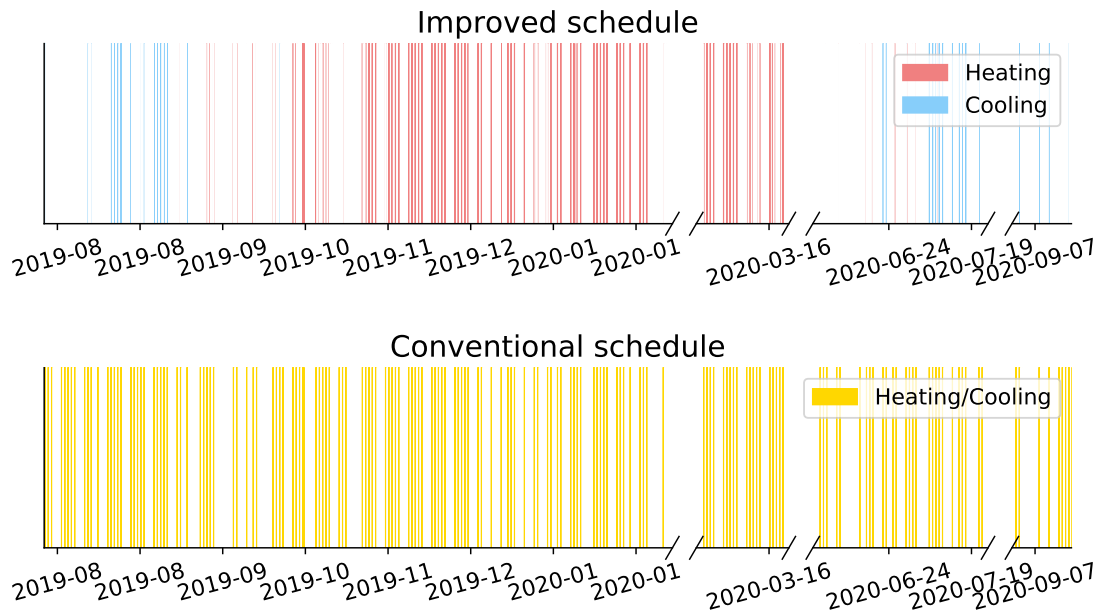


Figure 6.3 – Simulation 2, July 2019 - September 2020: comparison between conventional and improved ON/OFF schedules of the HVAC system (segment's width = ON or OFF duration).

As can be seen in Table 6.1, the share of spring days in the dataset is significantly lower than shares allocated to other seasons due to the effects of COVID-19. Regardless, some moderate potential energy savings due to reduced heating were obtained. The summer and autumn seasons contributed the most to potential energy savings, with relative daily shares of 0.81% and 0.71%, respectively. The winter season, instead, did not bring substantial potential energy savings, as it is characterized by higher occupancy rate and temperatures being frequently below the heating cut-off value. Particularly, the building's occupants spent 99.3% of their working time in winter indoors as opposed to 94.8% in summer. In autumn, the same metric showed 97.5% of indoor occupation time, suggesting that potential energy savings during this season occurred mostly due to mild October temperatures that fell inside the dead-band¹.

Table 6.1 – Distribution of potential energy savings among seasons

Season	Total days	Savings share	Daily savings
Winter	40	5.3%	0.13%
Spring	15	8.3%	0.55%
Summer	50	40.4%	0.81%
Autumn	63	45.2%	0.71%

¹A temperature range between T_{off}^{heat} and T_{off}^{cool}

The results of a similar analysis but related to heating and cooling functionalities of the HVAC system are presented in Table 6.2. The amount of heating and cooling days was calculated based on the occurrences of respective actions, resulting in 27 days during which the HVAC system was completely switched off. The values obtained for relative daily savings of two HVAC functionalities confirm the hypothesis stated earlier that cooling management results in higher potential energy savings in Portugal.

Table 6.2 – Distribution of potential energy savings among HVAC functionalities

Functionality	Total days	Savings share	Daily savings
Heating	106	44.1%	0.41%
Cooling	35	22.8%	0.65%

Table 6.3 provides estimations of respective financial savings that would be achieved for both simulations. Presented electricity prices correspond to Portuguese electricity tariffs established for non-household consumers in band-IB² and reflect the average value over the period [Eurostat 2021]. The resulting potential financial savings for Simulations 1 and 2 represent 2.6% and 4.4%, respectively, of the total budget spent on electricity. However, one should keep in mind that in the current case study, the share of HVAC consumption in the building's overall load demand equals 17%, which is much less than 44%, originally indicated in the statistics for commercial buildings [U.S. EIA 2016].

Table 6.3 – Evaluation of potential energy savings in a university building in Portugal

Simulation	Savings [%]	Savings [kWh]	Price [€/kWh]	Savings [€]
October 2018 - March 2019	15.4%	4491	0.1771	795
July 2019 - September 2020	25.7%	12935	0.1776	2297

6.6 Conclusion

To summarize, this chapter presented a novel occupancy-centric rule-based HVAC automation algorithm that generates day-ahead ON/OFF operational schedules for heating and cooling systems. Two case studies conducted in a university building in Portugal have demonstrated potential energy savings of 15.4% and 25.7% on 6-month and 15-month datasets. Moreover, a higher savings potential of cooling systems, as opposed to heating systems, was discovered. The proposed HVAC scheduling method serves as a promising example of building automation digital energy service that integrates seamlessly into the broader ICT-based platform and benefits largely from occupancy forecasting introduced in Chapter 5. Further considerations of deploying the building automation digital energy service should include the following:

²Band-IB corresponds to non-household consumers with an annual consumption in the range 20-500 MWh

- To precise the amount of potential energy savings and provide superior thermal comfort to buildings' occupants, one should consider the exact configuration and the number of respective HVAC units in possession. In the current case study, the whole-building level HVAC operations were considered, while in practice, the HVAC-setting in large buildings is more complicated. Thus, HVAC management should resort to higher resolution, such as zone-level or room-level.
- Being computationally lightweight due to its rule-based nature, the algorithm is suitable for online applications, while generated improved schedules can be efficiently deployed using BMS. However, according to experience, modifying standard algorithms contained in BMS is rather complicated. Thus, one should consider the technical capabilities of enabling automation in advance.
- To achieve the maximum value from building automation and predict the benefits with higher precision, the degree of occupants' potential impact on the building's operations should be assessed. Particularly, occupants' possibility of regulating temperature through thermostats or enforcing HVAC operations outside periods defined by the automation algorithm can largely affect potential energy savings.

The following Part II of the current thesis zooms out from the smart building framework to explore digitalization at the level of smart community. Particularly, it investigates the important role of electric mobility in facilitating the integration of utility-scale renewables and increasing the overall energy system's efficiency, flexibility, and reliability. Moreover, it demonstrates the benefits of intertwining energy infrastructure, buildings, and transportation.

Energy Management in Smart Communities

Part II

7 The Smart Community Framework

Zooming out from the smart building to the smart community level helps to understand further benefits of digitally interconnecting energy infrastructure, buildings, and transportation. Mobility electrification in particular represents a great example of sector coupling that contributes to energy efficiency, flexibility, and reliability of the overall energy system. This chapter introduces an ICT-based framework implemented on top of existing physical assets that aids to investigate the importance of intelligent green mobility in supporting digitalization of the energy industry.

Related **publications**:

[8] Marina Dorokhova, Jérémie Vianin, Jean-Marie Alder, Christophe Ballif, Nicolas Wyrsh and David Wannier. *A blockchain-supported framework for charging management of electric vehicles*. Submitted, 2021

7.1 Background and motivation

As demonstrated in Part I of the current thesis, buildings' digitalization represents a massive potential for increasing energy efficiency. Therefore, it comes without a surprise that smart buildings are often regarded as protagonists of the energy transition. However, when one considers them as an integral part of energy infrastructure, it becomes obvious that smart buildings' contribution to reaching sustainability goals goes far beyond their individual performance and characteristics. Thus, Part II of the current thesis aims to zoom out from the smart building framework and investigate the impact of digitalization phenomena on the broader energy ecosystem at the level of smart community:

Smart Community

An entity that leverages integration of infrastructure, data, and ICT to enhance energy efficiency, create cost savings, and deliver accessible services that improve the quality of life for its participants.

Thanks to smart buildings' inherent capability to connect and interact with other actors of the energy landscape, smart communities reap the benefits of using advanced technology to tightly intertwine energy infrastructure, buildings, and transportation. Often referred to as sector coupling [Ramsebner 2021], the idea of interconnecting energy-consuming sectors with the power sector aims to facilitate the integration of utility-scale RES and increase the efficiency, flexibility, and reliability of the overall energy system. Mobility electrification is regarded as an integral component of urban energy in the future due to Electric Vehicles' (EVs) inherent ability to act as distributed energy storage. Being parked 95% of their time [Barter 2013], EVs represent an untapped potential for implementing digital energy management solutions at the smart community level. Thus, the research conducted in Part II of the current thesis focuses on the important role of intelligent green mobility in supporting further digitalization of the electric power sector.

7.2 The Digitalization project

The Digitalization project is conducted within the research framework of the Swiss Competence Center for Energy Research (SCCER) FURIES mandated to develop, promote, and deploy power grid-related innovative solutions toward the implementation of the Swiss Energy Strategy 2050 [Innosuisse 2019]. The project, established as part of overall SCCER FURIES activities related to the Swiss National Action Plan on Digitalization [SERI 2020], aims to develop in-country ICT competencies in the various areas of the electrical infrastructure digitalization. In particular, it intends to explore the usage of digital platforms to coordinate action between energy players and achieve the following research objectives, which constitute a part of SCCER FURIES overarching goals in digitalization:

7.3. The ICT-based smart community framework

- Evaluate the potential of EV charging optimization on increasing the self-consumption of smart buildings with PV installations, providing peak-shaving capabilities to the grid, and reducing the impact of electric mobility on consumers' electrical bills
- Simplify the exchange of energy between smart buildings, EV charging stations, and EVs using digital methods based on blockchain technology
- Assess the opportunities and limitations of implementing digital energy services related to EV management in smart communities

The project's consortium is composed of more than 30 partners coming from both academia and industry. However, to achieve the specific objectives indicated in the bullet points above, EPFL has particularly collaborated with researchers from HES-SO Valais. The demonstration site, chosen to showcase the implementation of some of the project's digitalization activities, is located in Val d'Hérens alpine region in Switzerland. This site integrates a network of EVs and charging stations available to guests staying in the hotels of the region. Each of the 8 partner hotels owns one charging station and at least one EV, allowing their guests to explore the region with maximum independence and minimum harm to the environment. The EVs are rented to the hotels' guests daily free of charge, based on the principle "pay what you want". Thus, the guests themselves determine the appropriate level of compensation for this EV renting service, offered by the Green Mobility initiative in Val d'Hérens [Val d'Herens 2019].

7.3 The ICT-based smart community framework

Figure 7.1 depicts the architecture of the ICT-based smart community framework that brings together physical and digital assets. Particularly, in the proposed framework, the main components of the infrastructure are tightly intertwined with data management solutions and digital energy services to streamline the functioning of the smart community. The interconnections between the various elements of the ICT-based framework are depicted using the flow method, where information exchanges are shown using the dashed line, while the power exchanges are depicted with the solid line. One should note that the demonstrated architecture is scalable and can accommodate several smart buildings, EVs, and charging stations, even though they are drawn in single quantities. The following elements compose the ICT-based framework of the smart community:

- The **PV installation** belongs to the smart building and is usually located on its rooftop. The renewable power from PV is used to satisfy the smart building's load demand. The excess of the PV production is supplied to the charging station when needed or is fed back to the utility grid by the smart building.
- The **utility grid** provides power to the smart building and charging station whenever there is demand.

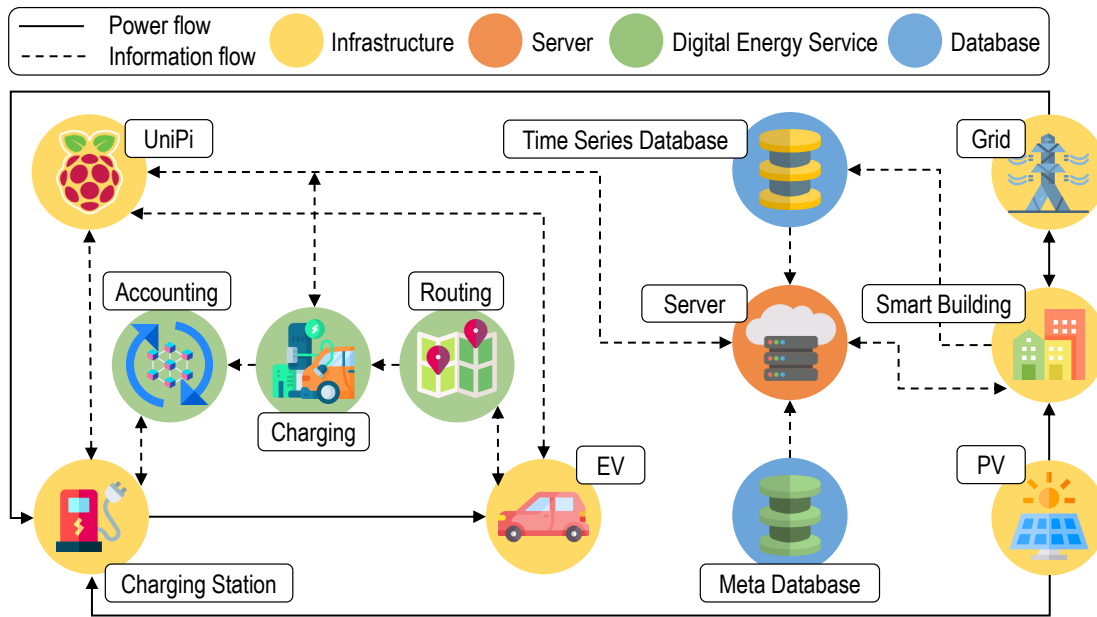


Figure 7.1 – The architecture of the ICT-based smart community framework

- The **smart building** is the main physical entity in the framework architecture. The smart building is characterized by its load demand, which can be satisfied either by the rooftop PV or by the utility grid. Moreover, the smart building owns the charging stations and at least one EV. The smart building sends the smart meter power measurements and the information about PV production to Time Series Database (TSDB). The building owner or manager interacts with the Server to issue charging requests to the charging station in manual mode.
- The **charging station** itself is not endowed with the layer of intelligence. Therefore all interactions with the charging station are conducted through the UniPi controller. The charging station is either AC or DC and can be operated in both manual and automatic modes. In manual mode, the maximum charging current is set by the smart building owner, while in automatic mode, the current is regulated according to the optimized EV charging strategy produced by the EV charging digital energy service presented in Chapter 9. The charging process is supported by both the utility grid and PV installation. Once the charging is complete, the charging station returns to UniPi the charging status and the amount of energy consumed in kWh.
- The **UniPi** is a programmable logic controller developed by HES-SO Valais and mounted inside the charging station. The UniPi enables the automatic control of the charging process through the commands received from the Server. The internal Application Programming Interface (API) allows the UniPi to send requests to the charging station using Modbus in write and read modes. Particularly, the UniPi can set the charging station's status, the energy consumption required, and the maximum amount of amps

the charger can deliver to the EV. Once the charging is complete, the UniPi returns the overall amount of energy consumed during the charging process in kWh to the Server.

- The **EV** on the scheme represents both the vehicle itself and the EV driver, who interacts with the UniPi of the charging station using a mobile application. When the EV arrives at the charging station, the driver optionally sends out the information about the EV's current State of Charge (SOC) and the amount of time the EV can spend at the charging station until the next departure. This information is used to optimize the charging process by means of the EV charging digital energy service proposed in Chapter 9. If the driver provides no supplementary information at arrival, the charging process proceeds without an optimization feature.
- The **Time Series Database (TSDB)** stores all data collected from the smart building. Besides the load consumption and PV production, the information might contain hot water usage and other measurements related to the building's equipment, such as heat pumps, boilers, energy storage, etc. TSDB feeds the data to the Server for visualization purposes in the front-end and for determining the optimal EV charging strategy.
- The **Meta Database** contains the smart building-related data used for creating the front-end of the visualization dashboard. The information includes the building's profile, the list and the characteristics of the installed equipment, etc. When the framework scales up to include multiple smart buildings, Meta Database becomes particularly useful for differentiating the static building-related data, such as installed equipment inventory, from dynamic time-series data, such as PV production, stored in TSDB.
- The **Server** links the elements of the framework architecture and enables communication from within. Moreover, it provides the visual dashboard to the smart building owner, where the latter can view the information in the interactive mode and issue the requests to the charging station. The data displayed include the building's measurements and the information related to the concluded charging processes such as charging start and finish times, duration, cost, and energy content. Particularly, the smart building owner can browse through the charging history of either the charging station or the EV that belongs to the building.
- The **EV routing** digital energy service helps the EV driver to navigate the roads of smart community in an energy-friendly way. Particularly, it resorts to Reinforcement Learning (RL) techniques to generate energy feasible paths from starting point to the destination. Notably, the EV routing digital service considers both the recharging possibilities on the way using intermediary charging stations and the ability of EVs to recuperate energy. The information about required charging stops is used by the EV driver to plan ahead for the charging sessions. Chapter 8 of the current thesis presents the details of the EV routing digital energy service.
- The **EV charging** digital energy service aims to control the EV charging process by means of RL with the objective to maximize the smart building's PV self-consumption while

satisfying the EV's charging needs and maximizing the SOC at departure. Resulting EV charging strategy is sent to the Server and transmitted by the UniPi to the charging station in the automatic control mode. More details on the design and implementation of the RL-based digital EV charging service are presented in Chapter 9 of this thesis.

- The **EV accounting** digital energy service is based on Ethereum blockchain technology and allows the EV and charging station owners to monitor the EV charging process while being securely credited for respective energy flows. The smart contract at the core of the digital service ensures the correctness of energy accounting for participating entities, while specifically designed web and mobile interfaces provide a user-friendly charging experience. Chapter 10 of the current thesis elaborates on the implementation of the blockchain-supported digital energy accounting service.

8 Routing of Electric Vehicles

The switch to electric mobility is regarded as one of the most promising solutions to combat climate change by decreasing greenhouse gas emissions. However, despite the rapid proliferation of Electric Vehicles (EVs) worldwide, their limited cruising range and sparse charging infrastructure hinder the real potential for the massive uptake of new-generation green transportation. To mitigate the drivers' range anxiety problem and successfully account for the specific features of EV technology, the EV routing method based on reinforcement learning is implemented.

The main **highlights** and **contributions** of this chapter are:

- The formulation of the EV-specific routing problem based on Markov decision process in the graph-theoretical context. The definition considers both recharging possibilities at intermediary charging stations and the ability of EVs to recuperate energy.
- The implementation of a model-free reinforcement learning algorithm that serves as an example of a possible method to solve the formulated EV routing problem with energy feasibility objective. A real-world case study demonstrated that 100% generated paths rendered energy feasible, while 92% of them represented near-optimal charging solutions.
- The results highlight the importance of extending the pool of routing objectives, considering additional characteristics of the charging process, and dedicating efforts to improve the scalability of the suggested methodology by choosing a suitable reinforcement learning algorithm.

Related **publications**:

[6] Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh. *Routing of electric vehicles with intermediary charging stations: A reinforcement learning approach*. Frontiers in Big Data, 4:586481, 2021. DOI: 10.3389/fdata.2021.586481

8.1 Background and motivation

The importance of EVs has increased significantly over the past few years, with growing concerns about climate change, volatile prices of fossil fuels, and energy dependencies between countries. Although the overall level of Greenhouse Gas (GHG) emissions has lately dropped considerably in the European Union (EU), the transport sector has not followed this trend. Particularly, it increased its relative share of GHG emissions to 27% in 2019, 72% of which were contributed by road transport [EEA 2020b]. Therefore, switching to electric mobility is seen as a primary mean of reaching emissions' reduction targets set by EU for 2030 [EC 2020].

Although the EV deployment grows fast around the world (+40% in 2019 compared to 2018), with Europe accounting for 24% of the global fleet [IEA 2020a], specific barriers for a massive uptake of EVs still exist. Researchers in [Noel 2020] identify technical, economic, social, and political barriers of broad adoption of EVs with limited cruising range and sparse charging infrastructure prevailing at present. These barriers are in the essence of the "range anxiety problem", defined as a fear that an EV will not have sufficient charge to reach its destination. However, optimal EV route planning together with longer-range EVs entering the market can mitigate this problem.

Route planning strategies have been widely researched for conventional vehicles powered by fossil fuels. However, solving the same problem for EVs requires a thorough consideration of EV-specific characteristics, such as limited battery capacity and ability to recuperate energy. Moreover, inadequate sparse charging infrastructure and long charging times call for selective choice of the charging station, additionally influenced by the price of electricity, expected charging power, distance from EV to the charging station, current State of Charge (SOC), expected waiting and charging times, and incentives from electricity providers. Another difficulty in route planning for EVs lies in the choice of the optimization goal. Conventional routing algorithms, such as Dijkstra [Dijkstra 1959], yield either the least traveled time or the least traveled distance. However, none of these options guarantees the energy feasibility of the generated route. Therefore, a need for EV-specific routing algorithms that strive for energy efficiency emerged.

The remainder of this chapter is organized as follows. Section 8.2 reviews the recent literature in the field of EV routing, especially focusing on the EV-specific features considered and the underlying methodology. Section 8.3 presents the proposed approach for solving the EV routing problem, particularly elaborating on the mathematical formulation and choice of the appropriate algorithm. Section 8.4 describes the case study conducted to validate the methodology. Section 8.5 focuses on the achieved results and discusses the advantages and limitations of the chosen algorithm. Section 8.6 concludes the study and suggests further improvements to the EV routing digital energy service.

8.2 State of the art

The algorithms in the field vary significantly by EV-specific features considered, complexity of the methodology, and application use cases. The first group of algorithms uses detailed energy consumption models that respect the ability of EVs to recuperate energy but neglect the possibility of recharging the battery on the way. Researchers in [Cauwer 2019] used the shortest path algorithm to find the optimal energy route on a weighted graph with a data-driven prediction of energy consumption. Authors in [Abousleiman 2014] deployed the ant colony and particle swarm optimization to generate the most energy-efficient route. Despite being fast, the solution is tedious to formulate and requires adaptation to different cases of EV usage. An interesting approach based on learning from historical driving data was demonstrated in [Bozorgi 2017]. The proposed solution aims to minimize both energy consumption and travel time while accommodating particular driving habits by considering variable speed profiles.

The second group of algorithms focuses on EV's interaction with charging stations while considering constant energy consumption without energy recuperation. [Sweda 2012] used approximate dynamic programming to minimize traveling and recharging costs. [Daanish 2017] deployed a nearest neighbor search-based algorithm to find the shortest energy-efficient path. Researchers in [Schoenberg 2019] and [Tang 2019] proposed algorithms to reduce the total travel time. The prior suggested a multi-criterion shortest path search with an adaptive charging strategy. The latter solved a joint routing and charging scheduling optimization problem that additionally minimizes the monetary cost.

The third group demonstrates an improvement in EV routing by considering both energy recuperation and battery recharging. An approach based on dynamic programming was proposed in [Pourazarm 2014] to minimize total travel time in the road network defined as a graph. Despite successful application for one car, the approach showed poor scalability in terms of convergence speed when the number of vehicles increased. [Morlock 2019] suggested a trip planner that solves a MILP to reduce the overall trip time. The authors introduced the driving speed as an additional degree of freedom and forecasted energy consumption from historical data. However, their approach works only along the desired route and does not consider alternative trajectories.

Although most of the proposed algorithms deal with route planning for casual EV driving, the efforts are made to adapt EVs for specific use cases of customer serving and delivery operations. Researchers in [Schneider 2014] deployed a hybrid heuristic search algorithm to minimize the total time that consists of travel time, recharging time, and time spent at the customer. Authors in [Mao 2020] aimed for the same goal with battery swapping and fast charging options using improved ant colony optimization. [Felipe 2014] used simulated annealing to find a feasible route while determining the amount of energy to be recharged at the charging station along with the type of charging technology. Despite considering the recharging possibilities on the way, these works neglect the ability of EVs to recuperate energy by assuming constant energy consumption proportional to the travel distance.

The last group collects very few works in the field that utilize Machine Learning (ML)-based approaches, particularly Reinforcement Learning (RL), to solve the EV-specific route planning problem. Previously, the success of using RL, namely the policy gradient algorithm, was demonstrated in [Nazari 2018] to minimize the total route length of a conventional fossil-fuel vehicle. Additionally, researchers in [Zhang 2020d] used actor-critic RL to minimize the energy consumption on the route without recharging opportunities. In [Zhang 2020a], a deep RL approach was proposed to reduce both travel time and distance while different charging modes and occupation of charging spots were considered. To address the highlighted drawbacks in the EV path planning and promote the usage of ML in the field, this thesis proposes a novel Markov Decision Process (MDP) formulation of the EV-specific routing problem, suitable to be resolved by RL. The suggested MDP operates in a graph-theoretical context and aims to generate energy feasible paths for EV from source to target. Moreover, it considers recharging possibilities on the way through intermediary charging stations and the ability of EVs to recuperate energy through the elevation profile of the road network. To demonstrate a possible method of resolving the suggested MDP, a model-free RL algorithm is deployed in the current chapter.

8.3 Methodology

RL is a powerful tool for decision-making and control that applies to nonlinear and stochastic systems. As shown in Figure 8.1, the concept of RL consists of an agent that interacts with the environment through actions similar to how a controller influences a technical system through control signals. To provide the modeling capabilities of the agent's behavior and simulate the agent's way of taking decisions, the problem has to be formulated as an MDP. A standard MDP (S, A, P, R) consists of the following elements: a finite set of states S , a finite set of actions A , state transition probability matrix P , and rewards function R . Therefore, the agent receives active feedback in the form of rewards when it changes the system's state by executing specific actions in the environment. In the context of the EV-specific routing problem, the agent represents an EV, and the environment mathematically describes a road network in a graph-theoretical setting.




Figure 8.1 – The concept of reinforcement learning

8.3.1 Environment

The road network can be modeled as a simple undirected weighted graph $G = (V, E)$ as follows:

- $V = \{1, \dots, n\}$ is the set of n nodes representing the points of interest on the map. The subset of these nodes $C = \{1, \dots, m\} \subset V$ can provide recharging capabilities to EVs. Each of the nodes $v_i \in V$ can serve both as a source v_0 and as a target v_f that are EV's starting and destination points, respectively. To consider the ability of EVs to recuperate energy when moving downhill, each node $v_i \in V$ is characterized by its elevation z_i .
- $E \subset \mathbb{R}$ is the set of weighted edges that connect the nodes on the graph. Each edge can be defined as an unordered pair $\{v_i, v_j\}$, where $v_i \neq v_j$. There are no multiple edges that are incident to the same two nodes. As the graph G is undirected, the edges are equivalent to two-way roads in the real world. The weights of the edges correspond to the energy costs required to traverse the edge.

The definition of edges' weights was adapted from [Bozorgi 2017]. Therefore, the energy cost E_{ij} between two nodes v_i and v_j can be determined as follows:

$$E_{ij} = E_{flat_{ij}} + E_{inclined_{ij}} + E_{other_{ij}} \quad (8.1)$$

Where $E_{flat_{ij}}$ and $E_{inclined_{ij}}$ represent EV's energy consumption on flat and inclined surfaces, respectively. The term $E_{other_{ij}}$ signifies additional energy costs that depend on road type, urbanization, weather conditions, and usage of auxiliary components [Li 2016]. For the sake of simplicity, $E_{other_{ij}}$ is neglected in the current work. The basic energy consumption on the flat road can be determined according to Equation 8.2, where h is the EV's specific energy consumption and d_{ij} is the distance between nodes. For different models of EVs, the value of h is determined experimentally on the basis of 100 km according to typical driving cycles, such as Worldwide Harmonised Light Vehicles Test Procedure (WLTP) [EEMA 2017].

$$E_{flat_{ij}} = d_{ij}h \quad (8.2)$$

The contribution of an inclined surface to the energy consumption of EV is proportional to potential energy and can be calculated as follows:

$$E_{inclined_{ij}} = mg\Delta z/\eta \quad (8.3)$$

Where m is the combined mass of EV and its payload, g is the acceleration of gravity, $\Delta z = z_j - z_i$ is the elevation difference between nodes, and η is the transmission efficiency of the EV. The value of $E_{inclined_{ij}}$ is responsible for EV's energy recuperation ability. In downhill, $\Delta z < 0$, therefore $E_{inclined_{ij}} < 0$ and EV can recuperate energy if $|E_{inclined_{ij}}| > E_{flat_{ij}}$. In contrast, $\Delta z > 0$ when EV moves uphill, thus $E_{inclined_{ij}} > 0$ and additional energy has to be spent. If two nodes have no edge connecting them, the weight $E_{ij} = \infty$, thus making it impossible for EV to traverse the graph in this direction.

8.3.2 Markov decision process

Figure 8.2 depicts a matrix representation of a combined state-action space, where the definition of states and actions is related to the graph-theoretical context of the problem.

		Actions							
		Set of nodes V				Set of charging stations C			
		(1,0)	...	(n,0)	(1,1)	...	(m,1)		
		(1, bat _{max})	(1, 0.9 · bat _{max})	...	(1, bat _{min})	...	(n, bat _{max})	...	(n, bat _{min})
States	(1, bat _{max})	inf				E_{ij}	inf		E_{ij}
	(1, 0.9 · bat _{max})	inf				E_{ij}	inf		E_{ij}
	...	inf				E_{ij}	inf		E_{ij}
	(1, bat _{min})	inf	inf	inf	inf	inf	inf	inf	inf
	...								
	(n, bat _{max})	E_{ij}				inf	E_{ij}		E_{ij}
	...	E_{ij}				inf	E_{ij}		E_{ij}
	(n, bat _{min})	inf	inf	inf	inf	inf	inf	inf	inf

Figure 8.2 – Matrix representation of a combined state-action space

State space S contains all possible states s that the agent can have when interacting with a given environment. For the case of EV routing, a state can be described as a vector $s = (location, charge)$, where $location \in V$ and $charge$ corresponds to the battery energy level. The latter is constrained due to battery's operational limits such as $bat_{min} \leq charge \leq bat_{max}$. The upper bound bat_{max} is imposed by the battery capacity and the lower bound bat_{min} is determined by the advised discharging policy. As most rechargeable batteries are not meant to be fully discharged, a minimum allowed SOC is set to avoid battery damage. In this research, it is assumed that $bat_{min} = 20\% bat_{max}$. Contrary to $location$, $charge$ is a continuous variable that might require discretization, depending on the selected RL algorithm. To achieve that, a uniform binning can be deployed, where the bin's lower bound determines the next state. The optimal number of bins can be identified experimentally, as a trade-off between the level of detail at which the problem is modeled and the size of the state space. With a large number of states and actions, the probability of visiting a particular state and performing a specific action decreases dramatically, thus deteriorating the performance, slowing down the training process, and exhibiting higher memory demands.

Action space A contains all possible actions that the agent can perform in the environment. An action can be described as a vector $a = (next_location, decision)$, where $next_location \in V$ and $decision$ indicates the charging intention at this location. If $next_location \in C$, the agent can choose whether to charge $decision = 1$ at this node or not $decision = 0$. If $next_location \notin C$, the agent has no choice and $decision = 0$. However, at any state s not all actions are available to the agent. The action a is considered available at state s only if $charge_s - E_{sa} \geq bat_{min}$, where E_{sa} is the energy cost to move from $location$ to $next_location$.

Reward function R is the most challenging part of the MDP formulation, which usually requires careful engineering. When interacting with the environment, the agent takes action from the current state, observes the new state, and receives feedback on the quality of its decision-making in the form of immediate reward R_t . Therefore, the agent continuously collects feedback from the environment over n steps required to complete the task. Ultimately, the agent learns the desired behavior by maximizing its cumulative discounted reward G_t , computed as follows:

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \gamma^3 R_{t+4} + \dots + \gamma^{n-1} R_{t+n} \quad (8.4)$$

Where discount factor γ that obeys $0 \leq \gamma \leq 1$ is used to emphasize the importance of the rewards achieved in the future. Therefore, one needs to find balance between considering immediate rewards only ($\gamma = 0$) and caring about distant future ($\gamma = 1$).

The shaping of the reward function R is particularly important as it should lead the agent towards achieving the objective in a fast and optimal way. Therefore, efficient reward functions facilitate the speed of convergence of RL algorithms and help the agent avoid being stuck in the local minima. The R can be mainly grouped into two categories: dense and sparse. Dense rewards are given almost at every state transition, while sparse rewards are infrequent. Dense rewards are more difficult to define as they require expertise and domain knowledge. Moreover, they are prone to limiting the agent's behavior from the discovery of creative solutions and facilitating strange behaviors learned unintentionally [Vecerik 2017]. Sparse rewards, instead, are easier to formulate, such as a binary signal indicating task completion or reaching the objective. However, they demand long training processes as most of the time, the agent does not receive any feedback from the environment. In the EV-specific routing problem, sparse rewards are selected. As the aim is to mainly incentivize only one type of behavior that ensures energy feasibility, the reward is set to 1 when reaching the target v_f from the source v_0 is done with charge level $charge \geq bat_{min}$. Rewarding the arrival to the final destination is essential for the agent's understanding that it has to explore the graph in a specific direction and not just wander around the environment. However, not all rewards have to be positive. Sometimes, rewards are used to penalize particular behaviors. In the current case of EV routing, the agent receives a negative reward equal to -1 when there are no available actions at the current state. In the real world, it would mean that EV has exhausted its battery capacity and thus got stuck on its route before reaching the destination.

In the current thesis, the state transition probability matrix P is not explicitly calculated due to the following assumptions in formulating the EV-specific routing problem. First, specific traffic conditions are not considered. It is common for drivers to plan their routes according to traffic congestion and even change them while driving. Therefore, the probability of choosing a particular road would need to be adjusted dynamically. Second, as the aim is to solve the routing problem for energy feasibility, the charging stations' occupation level and the time required for charging are not taken into account. Third, it is assumed that there are no partial recharges and that all EVs leave the charging station with the full battery. Moreover, although the behavior of the EV driver is presumed to be rational, in the real world, it is still stochastic. The drivers are free to choose the next points on their path according to any unforeseen events or their personal beliefs. To summarize, considering all the points discussed above, calculating the actual state transition probability matrix P that would accurately reflect real-world environment dynamics does not seem possible. Therefore, a model-free RL algorithm that operates regardless of any representation of P should be selected to solve the MDP of the EV-specific routing problem. To find the target policy that fully defines the agent's desired behavior, this work deploys the off-policy learning method that allows to do it independently from the followed exploratory policy.

8.3.3 Algorithm

As one of the possible methods to solve the suggested MDP formulation of the EV-specific routing problem, the Q-learning algorithm is chosen, which is a specific instance of temporal difference learning that looks only one step ahead. Moreover, it is suitable for discrete state and action spaces and is easily interpretable. The idea of Q-learning is to allow improvements for both target and exploratory policies. The target policy is a greedy policy that obeys the following definition:

$$\pi(s') = \underset{a'}{\operatorname{argmax}} Q(s', a') \quad (8.5)$$

Where π is the policy, Q is the action-value function, s' is the next state, and a' is some alternative action that maximizes the Q-value. The real behavioral policy that the agent follows is an ϵ -greedy policy that ensures continual exploration. The policy is defined as follows:

$$\pi(a|s) = \begin{cases} \epsilon/m + 1 - \epsilon, & \text{if } a^* = \underset{a \in A}{\operatorname{argmax}} Q(s, a) \\ \epsilon/m, & \text{otherwise} \end{cases} \quad (8.6)$$

Where s and a are the current state and action taken at this state, ϵ is the parameter that governs the exploration-exploitation trade-off, m is the number of actions available at the current state, and a^* is the best possible action. The Q-value function is updated according to

Bellman's optimality equation in the following way:

$$Q(s, a) \leftarrow Q(s, a) + \alpha [R(s, a) + \gamma \max_{a'} Q(s', a') - Q(s, a)] \quad (8.7)$$

Where $Q(s, a)$ is the Q-value of the current state and action pair, $R(s, a)$ is the observed reward after the action a is taken, and α is the learning rate bounded by $0 < \alpha \leq 1$. The latter determines to what extent the newly acquired information overrides the old information. The complete Q-learning algorithm is described in Figure 8.3.

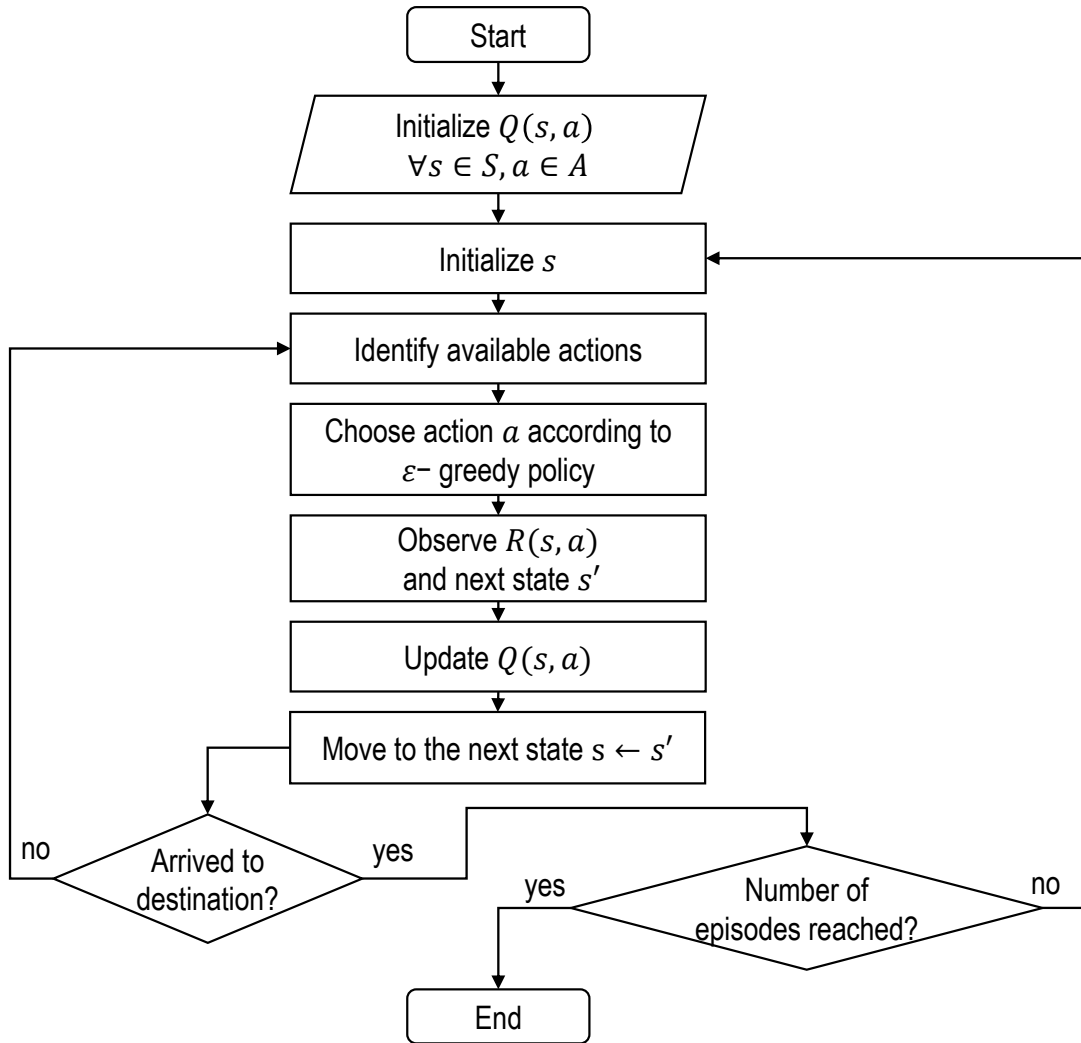


Figure 8.3 – One-step ahead Q-learning algorithm

Once the learning process, the duration of which is determined by the number of episodes, is completed, the resulting Q-table can be utilized for route planning purposes. Figure 8.4 depicts the decision-making process for constructing an energy feasible route from starting

point v_0 to destination v_f . Although commonly, most paths are started with the full battery equal to bat_{max} , an EV driver is free to input any value that corresponds to the EV's actual SOC. Afterward, the energy feasible trajectory leading to destination v_f is determined by the sequence of best actions chosen according to Equation 8.5. Notably, as the selected rewards scheme reinforces only the energy feasibility, the resulting route has to be refined to avoid a potentially excessive number of charging stops. Alternatively, nudging the agent to charge only when it is strictly necessary can be incorporated into the problem by reformulating the reward function R . The finalized trajectory that indicates the necessary stops at intermediary charging stations is then utilized to guide the EV driver from the starting point v_0 to destination v_f .

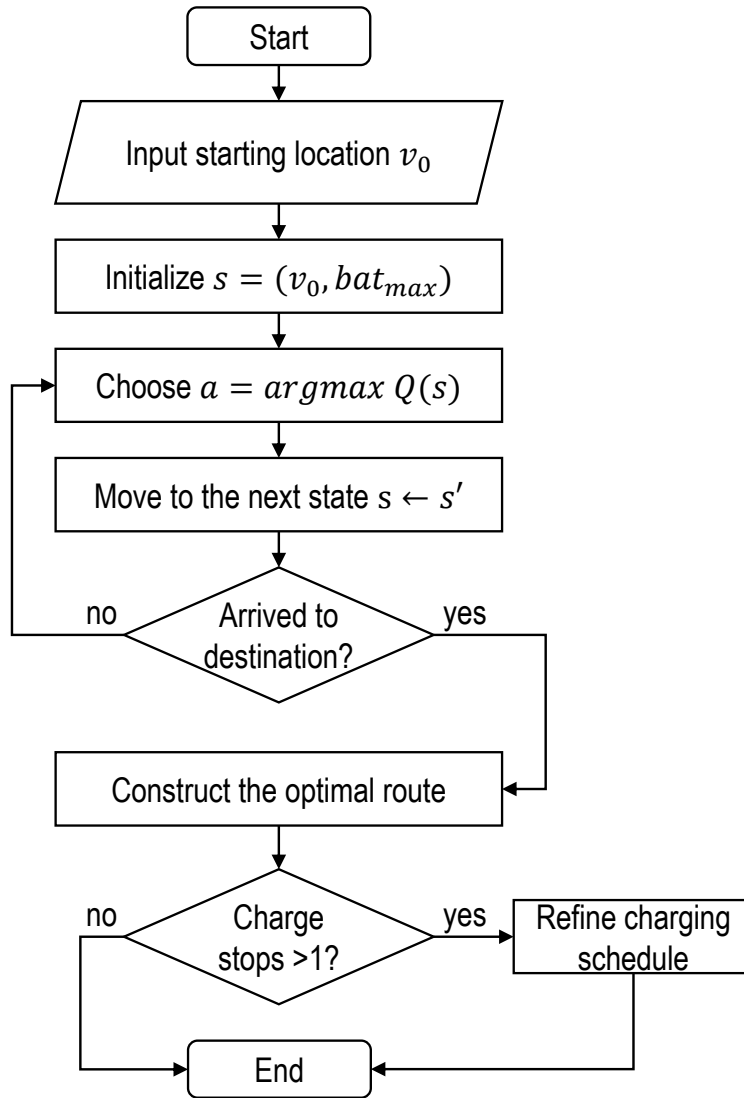


Figure 8.4 – Route planning process

8.4 Case study

To validate the suggested possible method for solving the EV-specific routing problem, a case study was created within the smart community framework proposed in Chapter 7. The case study deals with the section of the road network of Val d'Hérens alpine region in Switzerland, depicted in Figure 8.5. The undirected graph that describes the environment encompasses 66 nodes and 223 edges representing the points of interest and connection roads, respectively.

The landmarks comprise popular winter and summer destinations, prominent attractions and towns of the region, charging stations, and hotels that serve as the bases for EV rental. Notably, some nodes which are not specifically designated as charging stations in Figure 8.5 can also propose charging capabilities. The edges correspond to actual roads in the real world, extracted from Google Maps, with their thickness increasing with the nodes' relative remoteness. Each node is characterized by its geographical coordinates that consist of latitude, longitude, and elevation, equally extracted from Google Maps.

The agent is the EV defined by its battery capacity, energy consumption rate, and mass. In this case study, the Citroen C-Zero with 16 kWh battery and an average of 12.6 kWh energy consumption per 100 km [EVD 2021a] is used. The EV's mass is assumed to be equal to the gross vehicle weight of 1450 kg [EVD 2021a], defined as the maximum total safe weight counting vehicle's chassis, body, battery pack, electronics, accessories, driver, passengers, and cargo. The EV's transmission efficiency η is assumed to be equal to 0.85 for continuous variable transmission [Hofman 2010]. The continuous *charge* variable that determines the state space along with discrete *location* variable, is discretized into 50 bins uniformly.

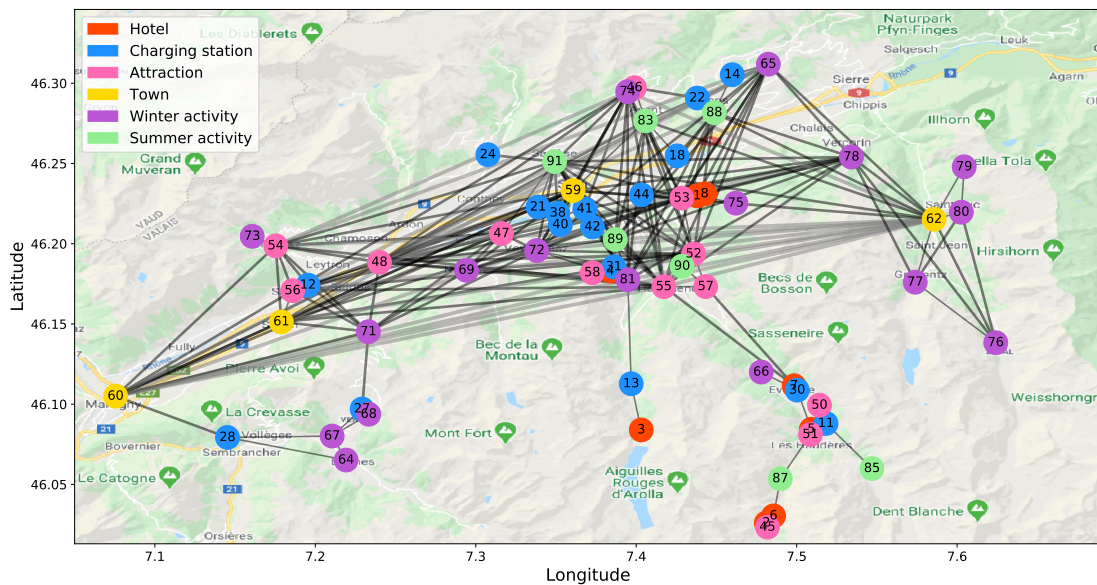


Figure 8.5 – Graph representation of the road network of the Val d'Hérens alpine region in Switzerland

8.5 Results

8.5.1 Training phase

The training procedure in RL is defined as a sequence of episodes, each representing the agent's journey along the path from source to target. Thus, an episode is considered complete when the agent reaches the target or has no available actions to perform due to a depleted battery. The number of episodes is determined experimentally and should be sufficient to achieve a stable Q-values matrix, initialized to zeros at the beginning of the training procedure. Such Q-matrix represents the maximum expected future rewards for each action at each state. The training convergence is achieved when the updates of the old Q-values become insignificant. Therefore, the agent learns the optimal policy once the algorithm converges. The parameters that govern the training process are set to the following values: discount factor $\gamma = 0.9$, learning rate $\alpha = 0.8$, and $\epsilon = 0.1$ in ϵ -greedy exploration policy. The values are tuned experimentally to ensure convergence and satisfactory execution speed.

Figure 8.6 depicts an example of the algorithm's learning curve, obtained during the training process with a fixed sample target. The x-axis denotes the number of episodes, while the y-axis represents the episode's training score determined as the mean of the scores obtained at each step of the episode. The step's score is calculated as the sum of all Q-values in the Q-matrix, normalized by the maximum observed Q-value. Therefore, the learning curve arrives at a plateau when the Q-matrix stabilizes. The Q-learning algorithm was programmed in Python, and the training procedure was executed on a personal laptop (Intel i7- 7600, 16 GB RAM).




Figure 8.6 – Example of a learning curve obtained during the training process. The bold line and the shaded region show the mean and the standard deviation of 5 runs.

In the current example, the algorithm converges after approximately 250'000 episodes, which takes 6.2 minutes. However, this reference execution time of the training procedure should not be regarded as absolute for two reasons. First, the duration of each episode varies according to the starting arbitrary state and exploratory actions taken by the agent. Therefore, a collection of training procedures with the same number of episodes would result in various training times and converge at a different number of episodes. Second, the chosen binning parameter has a large influence on the training duration. Particularly, a higher number of bins increases the size of the state space linearly, while the training's execution time is affected exponentially, as can be seen in Figure 8.7. Therefore, one should choose an appropriate value of the binning parameter based on a trade-off between the desired level of detail and available training time.

As mentioned previously, the training process uses a fixed target while the source is determined arbitrarily. Therefore, the algorithm requires re-training when a different destination is chosen. In a real-world application, such behavior would result in undesirable waiting times imposed on EV drivers. To avoid this, the EV routing algorithm should be pre-trained on the whole road network to obtain a collection of Q-matrices corresponding to different destinations, thus making its utilization convenient for users. Notably, any topological modifications of the selected road network, such as introducing additional nodes or removing existing ones, would equally require re-training the Q-learning algorithm. Therefore, one should investigate some alternative RL methods to resolve the suggested MDP that would clearly address the re-training limitation, thus enabling greater generalization to various road networks.




Figure 8.7 – The influence of binning parameter on the training time and state space's size. All experiments are conducted with a fixed number of episodes equal to 10000.

8.5.2 Validation phase

A series of experiments, where each node sequentially serves as a target, is carried out iteratively to test the consistency of the policy learned by the agent with the energy feasibility goal. Each experiment simulates an EV trip starting at a random node on the graph and finishing when the final destination is reached. For each target, the amount of experiments equals $N - 1$, where $N = 66$ is the total number of nodes in the selected road network. Thus, the overall number of experiments amounts to 4290. Besides verifying the EV's capability of arriving at the target without violating the bat_{min} constraint, the experiments are conducted to observe whether the EV stops to recharge only when it is strictly necessary. Although not accounted for in the reward function's design, excessive charging behavior is not preferable. Thus, observing the frequency of unnecessary charging stops is of interest for further improvements of the EV routing service. To simulate real conditions, it is assumed that the EV departs from the starting point with the battery fully charged.

The results demonstrate that 100% of generated routes are energy feasible, while 92% of them represent near-optimal charging decisions. The latter means that recharging schemes suggested by the algorithm comply with both following criteria. On the one hand, they give the agent a possibility to arrive at the destination, which is otherwise unreachable without charging. On the other hand, they neglect to charge when it is achievable to arrive at the destination without violating battery constraints. Moreover, the results show that in 80% of cases, the optimal number of charging stops was selected. Such a number was calculated using a verification procedure that analyzes the route with all possible combinations of the charging stations proposed by the algorithm. Although the method did not aim to optimize for the route length, an interesting observation occurred. In 83% of cases, the algorithm generated the shortest possible path when recharging was not required, as confirmed by the Dijkstra algorithm.

8.5.3 Discussion

The suggested MDP formulation of the EV-specific routing problem and the RL approach that was chosen as a possible way to solve it, both exhibit certain advantages and limitations in comparison to previous works in the field. On the bright side, the suggested method considers both main properties of EVs: the possibility to recharge on the way and the ability to recuperate energy. Although these features are crucial to model the agent's behavior close to real-world driving habits, taking into account both is uncommon, as shown in Section 8.2. Among the RL approaches particularly, the prior was neglected in [Zhang 2020d], while the latter was considered in [Zhang 2020a] through estimating energy consumption from rarely available historical data. Other advantages of the suggested methodology are threefold. First, as it was demonstrated in previous studies [Mocanu 2018] and will be shown in Chapter 9 experimentally, a trained RL agent requires less computing effort and less memory space than model-based techniques and MILP formulations of the EV routing problem such as

[Pourazarm 2014]. Thus, it can be deployed for online applications if successfully transferred to the real world. Second, formulating the problem in a graph setting and using the Q-learning algorithm that employs Q-matrix make results' interpretation more intuitive. Finally, the off-policy temporal difference continuously evaluates the returns from the environment and makes incremental updates using bootstrapping. Therefore, unlike the Monte-Carlo approach, it is unnecessary to wait until the episode terminates to judge the agent's behavior.

Despite the inherent advantages mentioned above, the suggested approach also has certain limitations influencing performance. First, Q-learning is suitable for problems with small to medium size of a state-action space as it stores information in the form of Q-tables. Once the dimensions of the problem increase, the algorithm scales poorly. In the proposed formulation, the state-action space's growth can come from expanding the road network or increasing the discretization parameter. To solve the scaling issue, one can use function approximators, such as neural networks or tile coding, or switch to policy-based RL. The second limitation comes from fixing the minimum required battery charge at the target to bat_{min} . As some destinations might not have charging stations, the EVs can get stuck without sufficient battery charge to start a new trip. Thus, one has to introduce an additional parameter bat_f that depends on v_f and ensures that the battery charge at the destination is sufficient to arrive at the closest charging station. The third limitation of the method's applicability is the need to retrain the algorithm when the destination is changed or any topological modifications occur to the road network. Thus, it should be clearly addressed to improve the method's convenience for end-users. Finally, the agent's evaluation on the same environment model used for training questions its real-world performance and the ability to handle stochastic perturbations.

8.6 Conclusion

To conclude, this chapter presented a novel mathematical formulation of the EV-specific routing problem in a graph-theoretical context and proposed a possible method to solve it using a model-free RL approach. Notably, the suggested methodology considers both recharging possibilities at intermediary charging stations and the ability of EVs to recuperate energy. Since the objective is to generate energy feasible paths, the proposed EV-routing solution can either constitute a standalone digital energy service or become an auxiliary instrument in broader EV charging management. In both cases, implementation of the following features would potentially increase the chances of EV routing adoption among the EV drivers:

- Although the route's energy feasibility is one of the primary concerns feeding the range anxiety problem, further development of EV technology combined with the increasing number of charging stations will render other objectives important. Thus, one should diversify the routing problem towards minimizing travel time, travel distance, total energy consumption, and the number of recharging stops. The proposed approach simplifies such diversification as it allows to tailor the desired goal by altering the rewards scheme while keeping the remainder of the problem formulation unchanged.

- To generate the routes that apply to real-world driving conditions and reflect the characteristics of regional charging infrastructure, one has to extend specific features of the charging process to include such parameters as charging time and intensity, price of electricity, charging slot availability, and recharging amount. Moreover, one can improve the EV routing digital service by including dynamic traffic conditions, thus affecting the actions' availability and the agent's energy consumption model. The latter can be advanced by accounting for the type of terrain, use of auxiliary loads, and weather.
- To ensure the EV routing algorithm's convenience for online usage and applicability to large road networks, one should devote efforts to improving the scalability of the suggested methodology. As shown in the case study, discretization of continuous variables significantly influences the size of the state-action space, thus impacting the algorithm's training time. Therefore, new RL approaches suitable for continuous state spaces have to be incorporated to resolve the defined EV routing MDP.

Although some RL fundamentals have been already presented in the current chapter, the following one constitutes a deeper dive into the topic of RL by transitioning towards deep Reinforcement Learning Control (RLC). Particularly, it focuses on developing the EV charging control method with the objective to simultaneously maximize PV self-consumption and EV's SOC at departure. Moreover, it addresses some of the limitations discussed in Section 8.5.3 and provides a thorough comparison of RLC with other popular control methods, such as rule-based control and model predictive control.

9 Energy Management of Electric Vehicles

In light of the pressing climate change issues, the inherent advantages of electric mobility will undeniably determine the future of transportation. However, the adoption of Electric Vehicles (EVs) worldwide requires significant energy management efforts to ensure their smooth integration into the broader energy system. To prevent potential grid operational threats from outweighing EVs' expected contribution to energy flexibility and greenhouse gas emissions reduction, the EV charging control algorithms resorting to reinforcement learning techniques are introduced.

The main **highlights** and **contributions** of this chapter are:

- The proposal of three different mathematical formulations in the form of Markov decision processes that enable reinforcement learning control with discrete, continuous, and parametrized types of action spaces for the energy system consisting of a smart building, Photovoltaic (PV) generation, EV, and utility grid.
- The deep reinforcement learning algorithms, namely Double Deep Q-Networks, Deep Deterministic Policy Gradient, and Parametrized Deep Q-Networks are applied to the problem of EV charging control with the aim to simultaneously maximize PV self-consumption and EV state of charge at departure.
- The results of benchmarking reinforcement learning algorithms against naive, rule-based and both deterministic and stochastic model predictive control highlight their suitability for online applications and superior fitness for future mobility systems due to faster execution times and reliance on growing volumes of data.

Related **publications**:

[7] Marina Dorokhova, Yann Martinson, Christophe Ballif and Nicolas Wyrsh. *Deep reinforcement learning control of electric vehicle charging in the presence of photovoltaic generation*. Submitted, 2021

9.1 Background and motivation

As discussed in Chapter 8, mobility electrification is destined to play an important role in alleviating the global climate change crisis due to its high potential for reducing GHG emissions. Although already today a typical EV produces just half of the GHG emitted during the life-cycle of an average fossil fuel-based car, the future impact is expected to be even more pronounced with a bigger share of renewables entering the electricity mix [Hall 2018]. Besides purely environmental benefits, EVs offer a range of other advantages, such as reduced cost and frequency of maintenance and a quieter, more enjoyable driving experience. Thus, coupled with governmental subsidies for EVs and a tendency to introduce more drastic measures, such as banning gasoline-fueled cars' sales by 2030 following the examples of Denmark and the UK, it comes without a surprise that more than 50 million EVs are expected to hit the roads worldwide in the next 5 years [Hoover 2021]. However, one has to note that the inherent advantages of green mobility do not come without additional challenges related to EVs' integration in the broader energy system.

Notably, increasing electrification of transport is expected to reshape the electricity load curve with the most pronounced effects for evening peak loads. Although EVs are unlikely to drastically drive up the overall electricity demand, the increase in peak loads can impose significant threats on the secure and stable operation of power systems due to capacity issues of grid infrastructures. Therefore, efficient control strategies are required to manage the charging processes of EVs to avoid significant grid investments and guarantee the stability of the electricity supply. Additionally, as driving patterns demonstrate, the EVs are parked approximately 95% of the time [Barter 2013], which gives the potential to intelligently shift the charging load, thus deploying smart energy management techniques. On the bright side of the increasing penetration of electric mobility is the opportunity to offer grid ancillary services to support the grid's various objectives. For example, using EVs can reduce energy costs, contribute to peak shaving, improve system balancing, and integrate a larger share of renewables into power production [Ding 2020]. However, the trade-off is to combine demand-response with seamless availability of EVs, such as the primary purpose of enabling mobility is served in a reliable and timely manner. Therefore, a need to develop EV charging control algorithms that ensure efficient integration of electric mobility into the grid has emerged.

9.1.1 Control methods

To effectively manage the charging processes of EVs, one has to choose between various control strategies. The three main broad classes of control methods include Rule-Based Control (RBC), Model Predictive Control (MPC), and Reinforcement Learning Control (RLC) [Kathirgamanathan 2021]. Indeed, there are certain advantages and disadvantages associated with each of these control techniques. Therefore, a suitable approach has to be chosen based on the trade-off between the complexity of the problem, control objectives, and available computational resources.

Rule-based control

RBC, previously utilized in Chapter 6, is the simplest control method that consists of a knowledge base and an inference engine. The prior defines the set of rules that govern the operations, and the latter takes actions based on the input data and the corresponding rules. Although RBC requires domain-specific expertise and knowledge to define the decision-making criteria, a rule-based nature is easy to understand as it provides transparent links between causes and effects. Therefore, solutions generated by RBC can be easily interpreted and justified. However, rule bases do not scale efficiently. Thus, RBC becomes inadequate for large and complex problems. In practice, solving complicated matters with RBC results in conflicting and overlapping rules that require significant human supervision. Moreover, RBC is not good at handling incomplete or incorrect data and can strangle creativity and knowledge discovery. Indeed, only the known rules can be applied, while detecting unusual relations between the elements of the system might be hindered by the RBC nature.

Model predictive control

MPC is a control approach based on solving a constrained optimization problem iteratively at each time step for a finite time horizon. The family of MPC methods includes deterministic and stochastic control, with the latter being able to handle uncertain disturbances without performance degradation and violation of constraints. The MPC consists of a system model, a set of constraints, an objective function, predictive trajectories of future disturbance inputs, and an optimization algorithm. The model-based nature of MPC poses significant demands to the quality and precision of the system's model while exhibiting limited generalization capabilities. Difficulties in handling nonlinear systems represent another major shortcoming of MPC. The method is not guaranteed to converge to a global optimum, nor is it guaranteed to provide a solution on time. Moreover, MPC is computationally complex, which makes it inappropriate for online use.

Reinforcement learning control

Taking into account the pros and cons of RBC and MPC approaches mentioned above, the current work aims to explore the application of RLC to the problem of EV charging control. Keeping RBC and MPC as benchmarking options, the RLC is selected for three main reasons. First, the fast-paced development of a connected world, thanks to the Internet of Things, Artificial Intelligence (AI), and digital transformation, coupled with increased penetration of EVs, provides conditions to generate an abundance of domain-specific data. Therefore, data-driven approaches thrive in such circumstances enabling better understanding and management of physical systems around us. However, one needs to keep in mind that RLC requires large volumes of training data, which are not always available or easy to collect from specific systems. Moreover, training demands high computational resources and long times, although a trained agent is fast and can be easily deployed online. The second reason for exploring RLC

lies in the availability of model-free algorithms, which can be characterized with better applicability to systems with unknown dynamics or affected by significant uncertainties. Moreover, they permit flexible configurations of environments, simplifying problem adaptation. Finally, the extension of RLC towards deep RLC allows scalability of decision-making problems that were previously intractable. For these reasons, the following literature review focuses explicitly on the RLC methods that one can apply to the energy management of battery-powered EVs. Notably, charging of hybrid EVs is out of the scope in the current work.

The remainder of this chapter is structured as follows. Section 9.2 presents the state of the art in the field of RLC. Section 9.3 describes the suggested methodology to solve the EV charging control problem, particularly focusing on formulating the MDP and selecting the appropriate RL algorithms. Section 9.4 provides details about the real-world dataset used for validating the proposed approach and presents benchmark control methods deployed for assessing the performance along with explanations of the evaluation procedure. Section 9.5 summarizes the results of extensive comparison across several control methods and discusses the applicability of the suggested RLC approach. Finally, Section 9.6 draws conclusions about the viability of the RL-supported EV charging management solution and suggests improvements to increase the adoption level of such digital energy service in the future.

9.2 State of the art

RLC applications to energy systems is a relatively new, though fast progressing field with the majority of works completed over the five-year preceding timeframe. Researchers in [Vázquez-Canteli 2019] conducted an extensive review of demand-response enabled by Reinforcement Learning (RL), while authors in [Mason 2019] looked at building energy management through the prism of RL algorithms. Both of these reviews included research on EV-related energy exchanges. The goal of this literature study is to highlight contributions that demonstrate a variety of objective functions, MDP formulations, and RL methods applied to the EV charging control problem. Some of the research pieces selected do not include EVs but deploy battery storage. Although availability is the crucial difference between the two, one can reuse the methodology for EV management. All works are divided into two main categories according to the action space used: discrete or continuous. No studies were found that deal with parametrized action spaces, which are hybrid between discrete and continuous.

Discrete action spaces are characterized by a finite number of available actions that an agent can choose from. Researchers in [Guan 2015] looked at the Energy Storage System (ESS) and Photovoltaic (PV) generation through Temporal Difference (TD) learning to minimize the electricity bill of residential customers. The authors in [Lee 2019] applied the Q-learning algorithm, an instance of $TD(\lambda)$, to the energy system composed of PV, a utility grid, an ESS, a single home, and controllable home appliances. The RLC delivered a 14% reduction of electricity bill compared to the optimization approach. The same methodology was implemented in [Kim 2018] with the additional inclusion of EVs and the Vehicle-to-Grid (V2G)

concept. Another attempt at Q-learning was demonstrated in [Wan 2018a] to determine cost-efficient EV charging schedules with an emphasis on the battery degradation cost. However, no renewable energy sources were taken into account. Researchers in [Chiş 2016] deployed fitted Q-iteration batch RL to reduce the long-term cost of EV charging. Although they did not consider renewable generation, an effort was made to predict electricity prices using an Artificial Neural Network (ANN). The action was defined by the amount of energy charged in the battery and was discretized into several equal levels ranging from 0 to battery capacity. The same algorithm in a similar problem setting was used by [Sadeghianpourhamami 2019] with the demand-response goal of load flattening. The suggested MDP formulation is scalable to any number of charging stations. The authors in [Lee 2020] combined the peak-shaving and the cost minimization objectives while implementing a Deep Q-Network (DQN) approach. Researchers in [Mocanu 2018] pursued the same goal for a PV, EV, and building appliances energy system using DQN and deterministic policy gradient algorithms. In [Di Wu 2018] both neural fitted Q-iteration and DQN were deployed to reduce long-term operating costs of a home energy management system coupled with ESS, PV, a utility grid, and an EV.

In a continuous action space, an action is some real-valued number usually bounded by the defined range. Often, such actions exist in multiple dimensions. The majority of the studies apply a Deep Deterministic Policy Gradient (DDPG), the most well-known RL technique to solve MDPs with continuous actions. The authors in [Zhang 2020b] used DDPG for EV charging control to maximize user's satisfaction with charging requirements while minimizing the charging expense. Specific grid-level constraints, such as power flow, voltage region, and capacity limits of the equipment, were examined in [Ding 2020] to maximize the profit of a distribution system operator in a scheme of multiple EVs. An ESS with home loads was studied in [Wan 2018b] to reduce total electricity costs through the actor-critic RL. However, none of these works considered renewable generation. In [Yu 2019], no EVs were involved in a system of ESS and PV generation under the energy cost minimization objective. Researchers in [Ye 2020] proposed an extension of DDPG with a prioritized experience replay. Their energy system model included a utility grid, PV, ESS, heat pump, boiler, and thermal energy storage. An approach using proximal policy optimization was demonstrated in [Li 2020] to deal with critical, shiftable, and controllable appliances, where EVs were enclosed in the latter category.

The works mentioned above were concerned with a single-agent RL, while few studies investigate the application of Multi-Agent Reinforcement Learning (MARL) approaches. The energy system in a multi-agent environment is represented by multiple actors that compete, cooperate or do both towards achieving a specific goal. Therefore, MARL finds application in the game-theoretical context. The authors in [Shin 2020b] considered PV, ESS, and EV charging stations to compute optimal charging schedules in a distributed manner. To reduce operational costs, they applied the CommNet algorithm on a discrete set of actions. An equilibrium-based MARL algorithm was used in [Fang 2019] to determine the energy charging and discharging scheduling of a residential microgrid composed of multiple EVs and both solar and wind generation. The double objective of increasing average profits and reaching maximum autonomy from the utility grid was framed as a Markov game. The authors

in [Claessens 2013] combined batch RL with function approximators and a market-based multi-agent system in a model-free manner. The system comprised of EVs and PV was studied under peak shaving and valley filling objectives. The results demonstrated a 50% reduction in grid's peak power. Although MARL is attractive for demand-response problems due to its distributed nature, experience sharing possibilities between agents, and potential speedup of learning using parallel computing, certain shortcomings might limit its application. Despite the curse of dimensionality problem and exploration-exploitation trade-off, which are also common for single-agent algorithms, the biggest challenge in MARL arises from the difficulty of specifying a learning goal. Moreover, the simultaneous learning of multiple agents leads to nonstationarity and an increased need for coordination between actors [Busoniu 2010].

The conducted literature review has shown that single-agent approaches prevail in EV charging control problems and that albeit the high diversity of considered energy systems, which often include EVs, building loads, PV, a utility grid, and ESS, the diversification in methodology and objectives is feeble. Moreover, few works use the advent of deep RL despite its success in robotics, transportation, and healthcare. Therefore, the aim of this research is three-fold:

- To fill the gap in the literature by demonstrating how the same EV charging problem can be formulated in different ways mathematically, thus triggering the application and the comparison of a variety of deep RLC algorithms
- To extend the EV control beyond the cost minimization objective
- To show how the performance of RLC compares with other control methods such as RBC and MPC

To fulfill these goals, a simple energy system composed of a utility grid, building load, PV generation, and a single EV is considered.

9.3 Methodology

9.3.1 Environment



Figure 9.1 – Energy system representation

The energy system in consideration consists of the following elements shown in Figure 9.1. A utility grid and PV generation represent energy sources, while an EV and a building constitute energy demands. One has to note that in this work, the V2G concept is not considered and EV is exploited as the only controllable load. In this context, the aim of EV charging control is to reach two overarching objectives: to increase PV self-consumption of the energy system and to achieve as high State of Charge (SOC) at the EV departure as possible.

Objectives and constraints

Let PV_t denote the PV production at time t , while L_t and EV_t are the building load demand and the EV charging demand at time t , respectively. The L_t can be supplied by both the utility grid and PV, resulting in corresponding quantities of L_t^G and L_t^{PV} at time t , thus:

$$L_t = L_t^G + L_t^{PV} \quad (9.1)$$

On the contrary, according to the modeling choice, the concurrent supply to the EV from multiple sources is forbidden by the charging station:

$$EV_t^G EV_t^{PV} = 0 \quad (9.2)$$

Where EV_t^G and EV_t^{PV} represent the power supplied to EV from the grid and PV, respectively. Following the definition in [Luthander 2015], the PV self-consumption objective can be formulated as follows:

$$\text{max Self-consumption} = \frac{\sum_{t=t_s}^{t_e} (L_t^{PV} + EV_t^{PV})}{\sum_{t=t_s}^{t_e} PV_t} \quad (9.3)$$

Where the nominator is the self-consumed part of PV, while the denominator is the total PV production. One has to note that t_s and t_e are used to denote the corresponding start and end of the time period, which in RLC is equivalent to the length of an episode. In this problem formulation, as EV is the only controllable load, an episode starts at the arrival of EV to the charging station at t_{arr} and ends with its departure at t_{dep} . Therefore, $t_s = t_{arr}$ and $t_e = t_{dep}$. The time resolution Δt chosen is hourly. According to proposed notations, the maximization of the SOC at departure objective is formulated in the following way:

$$\text{max } SOC_t, \text{ where } t = t_{dep} \quad (9.4)$$

In addition to building and EV simultaneity of power supply constraints, the following constraints are formulated that are necessary to respect for the EV charging control problem:

$$SOC_{min} \leq SOC_t \leq SOC_{max} \quad (9.5)$$

$$0 \leq EV_t^G, EV_t^{PV} \leq P_{max} \quad (9.6)$$

$$SOC_{t+1} = SOC_t + \eta^{EV} \frac{(EV_t^G + EV_t^{PV})\Delta t}{C_{bat}} \quad (9.7)$$

Where Equation 9.5 bounds the EV SOC. The upper bound $SOC_{max} = 1$ is imposed by the battery capacity, and the lower bound SOC_{min} is determined by the advised discharging policy. As most rechargeable batteries are not meant to be fully discharged for lifetime reasons, a minimum allowed SOC is set to avoid battery damage. In this research it is assumed that $SOC_{min} = 0.2$, following [Marra 2012]. One has to note that SOC_t denotes the SOC at the beginning of time t . Equation 9.6 imposes the upper limit on the EV charging power P_{max} , which is conditioned by the type of the charging station. Equation 9.7 determines the continuity of the SOC within an episode, where η^{EV} is the efficiency of the EV charging process, and C_{bat} is the EV battery capacity in Wh.

9.3.2 Markov decision process

To solve the EV charging problem using RLC, one has to formulate it according to MDP formalism, thus define S, A, P , and R . Similarly to Chapter 8, the state transition probability matrix P is not defined as it would be equivalent to modeling the full dynamics of the environment. Thus, the formulation represents incomplete knowledge, where the next state s' is determined by the environment depending on the current state s and the action a taken.

State space S represents a collection of all possible states s that an agent can have when interacting with a given environment. The system state s at time t is continuous and is defined as a vector $s_t = (PV_t, L_t, D_t, SOC_t)$, where in addition to PV generation, building load demand, and EV SOC, the parameter D_t denotes the time to EV departure from the charging station.

Action space A includes all possible actions a that an agent can perform in the environment. The definition of A determines the appropriate RL algorithm to solve the MDP. In the case of EV charging control problem, the three different ways to formulate A are suggested:

Discrete action space

Given the state s_t , action a_t represents EV's charging power bounded by Equation 9.6. Moreover, a_t reflects the power source, either grid or PV. Thus, $a_t \in \{0, \Delta e_g, \dots, \Delta e_g^{P_{max}}, \Delta e_{PV}, \dots, \Delta e_{PV}^{P_{max}}\}$, where both grid and PV power supplies are discretized according to Δe_g and Δe_{PV} . In total, 21 actions are defined, using the uniform binning procedure described previously in Chapter 8.

Continuous action space

The action a_t is specified as a real-valued number within $[-1, 1]$ bounds. The negative part $[-1, 0)$ corresponds to the power supply from the grid, while the positive part $(0, 1]$ is equivalent to the power supply from PV. To obtain the quantity of power supply, one has to multiply the action a_t by the maximum allowed power P_{max} from the charging station.

Parametrized action space

Figure 9.2 depicts the definition of a parametrized, discrete-continuous hybrid action space A . The action a_t is constructed as a tuple. The discrete part of the action specifies the source, either PV or grid, while the parameter part defines the amount of power supply within a $[0, 1]$ continuous range, identical for both sources. The value 1 of a parameter is equivalent to the power supply of P_{max} .

For all types of action spaces proposed, an additional PV production constraint has to be satisfied:

$$EV_t^{PV} \leq \max(PV_t - L_t, 0) \quad (9.8)$$

Where the right-hand side of Equation 9.8 represents the PV surplus after supplying the building load. Thus, the priority of the PV power supply is given to satisfy the L_t , while the EV charging from a renewable source comes secondarily.

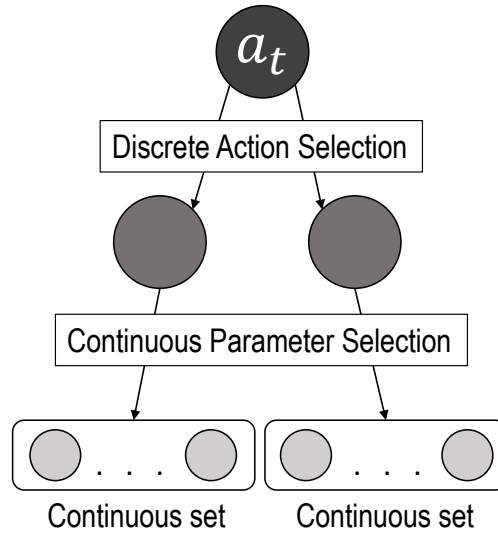


Figure 9.2 – Representation of parametrized action space

While in Chapter 8 the **reward** function R was used to guide the agent towards achieving a single objective, reward engineering for EV charging application requires an extension to multi-objective rewards. Therefore, the multi-objective sparse reward function R that helps the agent to maximize PV self-consumption while achieving the highest SOC possible at the EV's departure is defined using the following procedure.

To start, a notion of a maximum PV self-consumption SC_{max} achievable during an episode is computed as follows:

$$SC_{max} = \frac{L_{max}^{PV} + EV_{max}^{PV}}{\sum_{t=t_s}^{t_e} PV_t} \quad (9.9)$$

$$L_{max}^{PV} = \sum_{t=t_s}^{t_e} \min(PV_t, L_t) \quad (9.10)$$

$$EV_{max}^{PV} = \min\left(\sum_{t=t_s}^{t_e} \max(PV_t - L_t, 0), [SOC_{max} - SOC_{t_{arr}}]C_{bat}\right) \quad (9.11)$$

Therefore, Equation 9.11 is essentially a trade-off between the available PV surplus after satisfying building's demands and the available charge capacity in the EV battery. Indeed, the value of SC_{max} can only be calculated retrospectively, as the RLC agent has no knowledge about future values of PV_t and L_t . Therefore, the R function should provide the rewards at the end of the episode.

To design the multi-objective R , the inspiration was taken from challenging robotics environments [Plappert 2018]. The goal is 2-dimensional and describes the desired PV self-consumption and SOC at the end of the episode in the form of a tuple. Therefore, the reward function R can be defined as follows:

$$R = \begin{cases} 1, & \text{if } (SC_{max} - SC_{t_{dep}}) \leq \epsilon_{PV} \\ & \text{and } (SOC_{max} - SOC_{t_{dep}}) \leq \epsilon_{SOC} \\ 0, & \text{otherwise} \end{cases} \quad (9.12)$$

Where ϵ_{SOC} and ϵ_{PV} are the desired tolerance values and $SC_{t_{dep}}$ is the PV self-consumption achieved at t_{dep} . One has to note that such a formulation of R equally prioritizes the PV self-consumption and SOC at departure. In practice, however, one would give higher importance to SOC at departure as it directly impacts the EV's usability, while PV self-consumption only results in financial benefits. The sparse reward function R alleviates the burden of tuning the weights in case of combining multiple objectives using a dense function. Moreover, it allows defining a unique expression for calculating rewards, which is interchangeable between discrete, continuous, and parametrized environments. On the downside, sparse R requires specific techniques to speed up the training process and aid the exploration, especially when multiple objectives are combined. The Hindsight Experience Replay (HER) [Andrychowicz 2017] is introduced in the following subsection as a method that makes learning possible even if the rewards are sparse and binary.

To summarize, the EV charging control problem is defined as an MDP with the following characteristics:

- continuous state space S
- discrete, continuous or parametrized action space A
- sparse reward function R

9.3.3 Algorithms

Following the definition of the EV charging management problem in the form of MDP, one has to determine suitable algorithms to resolve it. As the current thesis aims to use RL techniques, their taxonomy, presented in Appendix A.4 and adopted from [Zhang 2020c], has to be considered to make an appropriate choice. Moreover, other characteristics of the problem at hand, particularly the type of state and action spaces, together with available time and computing resources, influence the selection. Thus, the following decision-making process has to be pursued to identify suited model-free RL algorithms:

- First, one has to choose the *policy-based* or the *value-based* class of the algorithms. The difference between the two is that the prior explicitly builds a policy representation and updates it iteratively during learning until the optimal policy is achieved. Thus, the policy defines the agent's behavior by mapping the states to actions. Instead, the latter class, as it was shown in Chapter 8 on the Q-learning example, stores only the value function that evaluates the goodness of being either in a certain state or taking a certain action in that state. Thus, the optimal policy can be derived from the learned value function by choosing the action with the best value. To bridge the two classes, one can choose actor-critic algorithms that efficiently deploy both *policy-based* and *value-based* approaches.
- Second, the *on-policy* or the *off-policy* learning method has to be selected. The prior aims to learn the exact policy the agent is executing, while the latter does it independently of the agents' actions. As it was already demonstrated in Chapter 8 on the Q-learning example, the values of the optimal policy are learned while the exploratory one is followed. The advantage of *off-policy* over *on-policy* methods is their higher sample efficiency, which, however, comes at the expense of lower stability when interacting with function approximators [Nachum 2017].
- Finally, the distinction between discrete and continuous state and action spaces has to be made. The current MDP is characterized by continuous state space, thus Artificial Neural Networks (ANNs) are used as function approximators to save computational time and memory resources. Instead, the action spaces vary, where continuous and parametrized ones help to preserve the high level of detail, otherwise downgraded through discretization.

Table 9.1 reflects on discussed above characteristics of chosen RLC algorithms.

Table 9.1 – Characteristics of chosen RL algorithms

Algorithm	Type	Learning	Discrete	Continuous	Hybrid
DDQN	value-based	off-policy	✓		
DDPG	value and policy-based	off-policy		✓	
P-DQN	value and policy-based	off-policy			✓

DDQN for discrete action space

The DDQN value-based algorithm [Van Hasselt 2015] is meant to solve the Q-values over-estimation problem of the DQN [Hasselt 2010] by decoupling action selection from action evaluation using two separate estimators, the online network and the target network. Algorithm 1 shows the pseudo-code of the DDQN implementation, where the training procedure is iterated over a specified number of episodes, each lasting from t_{arr} to t_{dep} of the EV.

Algorithm 1 Double Deep Q-Networks Learning (DDQN)

```

1: Initialize: online network  $Q_\theta$  and replay buffer  $\mathcal{R}$ ,
   target network  $Q_{\theta'}$  with weights  $\theta' \leftarrow \theta$ 
2: for each episode do
3:   observe current state  $s_t$ 
4:   for each step in the environment do
5:     select action  $a_t \sim \pi(Q_\theta(s_t))$  according to policy  $\pi$ 
6:     execute action  $a_t$ 
7:     observe next state  $s_{t+1}$  and reward  $r_t = R(s_t, a_t)$ 
8:     store  $(s_t, a_t, s_{t+1}, r_t)$  in replay buffer  $\mathcal{R}$ 
9:     update current state  $s_t \leftarrow s_{t+1}$ 
10:  end for
11:  for each update step do
12:    sample minibatch  $N$  of experiences:
        $e_i = (s_i, a_i, s_{i+1}, r_i)$  from replay buffer  $\mathcal{R}$ 
13:    compute expected Q-values:
        $Q^*(s_i, a_i) \approx r_i + \gamma \max_{a'} Q_{\theta'}(s_{i+1}, a')$ 
14:    compute loss  $L = \frac{1}{N} \sum_i (Q^*(s_i, a_i) - Q_\theta(s_i, a_i))^2$ 
15:    perform stochastic gradient descent step on  $L$ 
16:    update target network parameters:
        $\theta' \leftarrow \tau\theta + (1 - \tau)\theta'$ 
17:  end for
18: end for

```

At the beginning of the algorithm, one needs to create a replay buffer and to initialize online and target ANNs, where the target network is instantiated as a copy of the online network with the same weights. The first for-loop, lines 4-10, represents the interaction of the agent

with the environment, where the action a_t is selected using the exploration policy π based on the current state s_t . The DDQN algorithm uses an ϵ -greedy exploration policy, where the ϵ parameter decays to some value ϵ_{min} either linearly or exponentially. The second for-loop, lines 11-17, describes the learning procedure, enabled by the experience replay concept. The replay buffer \mathcal{R} serves as memory storage that holds all transactions between the agent and the environment in the form of tuples (s_t, a_t, s_{t+1}, r_t) , thus allowing the agent to reuse accumulated experience for the sake of better data efficiency. When the \mathcal{R} is full, the oldest experience is deleted to make space for the new one. Every n steps, the minibatch N of past transactions is sampled randomly from \mathcal{R} , thus breaking temporal relations between experiences and improving the stability of training. The learning process is three-fold. First, the expected Q-value $Q^*(s_i, a_i)$, line 13, is calculated using the Bellman equation based on the action chosen and evaluated by the online network Q_θ and the target network $Q_{\theta'}$, respectively. Second, the mean squared error loss is calculated between an expected $Q^*(s_i, a_i)$ and an actual $Q(s_i, a_i)$ Q-values and a gradient descent step is performed using Adam optimizer with selected learning rate α . Last, the parameters of the online network are slowly copied to the parameters of the target network using the rate of averaging $\tau \in (0, 1]$. In the case of hard-copying, the hyperparameter $\tau = 1$ and the weights' update is conducted every n steps.

Table 9.2 summarizes the choice of hyperparameters compared to the values of the original paper [Van Hasselt 2015]. These parameters are used in all DDQN experiments and are tuned empirically. To help the agent explore the environment, an ϵ -greedy policy is used with exponential decay from 0.9 to 0.1. The ANN configuration has two hidden layers of 32 nodes each with Rectified Linear Unit (ReLU) activations. The output layer size is equal to the size of the action space, and the activation is a softmax.

Table 9.2 – Hyperparameters of the DDQN algorithm

Hyperparameter	Symbol	Original value	Our value
Replay buffer size	\mathcal{R}	1×10^6	5×10^5
Minibatch size	N	32	32
Discount factor	γ	0.99	0.99
Averaging rate	τ	1	1
Learning rate	α	0.00025	0.001
Update every	n	1000	32

DDPG for continuous action space

The DDPG is a model-free off-policy algorithm for learning continuous actions [Lillicrap 2015]. Contrary to DQN-like algorithms, DDPG can handle continuous action spaces without discretization, thus alleviating the curse of dimensionality problem. As seen in the pseudo-code of the Algorithm 2, DDPG has an actor-critic architecture. The actor network $\mu(s|\theta^\mu)$ is the policy network that maps states to actions in a deterministic way. Thus it proposes the action to be executed. The critic network $Q(s, a|\theta^Q)$, instead, is the value network that judges the

quality of the state-action pair and thus evaluates the policy. Both networks are initialized randomly. The target critic Q' and the target actor μ' networks, instead, are created as copies of original networks, which can be seen in line 1. The purpose of target networks is to add stability to the training process.

Algorithm 2 Deep Deterministic Policy Gradient (DDPG)

```

1: Initialize: critic network  $Q(s, a|\theta^Q)$  and actor network  $\mu(s|\theta^\mu)$ , target critic network  $Q'$  and
   target actor network  $\mu'$  with weights  $\theta^{Q'} \leftarrow \theta^Q, \theta^{\mu'} \leftarrow \theta^\mu$ , replay buffer  $\mathcal{R}$ 
2: for each episode do
3:   observe current state  $s_t$ 
4:   initialize random process  $\mathcal{N}$  for action exploration
5:   for each step in the environment do
6:     select action  $a_t \sim \mu(s_t|\theta^\mu) + \mathcal{N}_t$ 
7:     execute action  $a_t$ 
8:     observe next state  $s_{t+1}$  and reward  $r_t = R(s_t, a_t)$ 
9:     store  $(s_t, a_t, s_{t+1}, r_t)$  in replay buffer  $\mathcal{R}$ 
10:    update current state  $s_t \leftarrow s_{t+1}$ 
11:    sample minibatch  $N$  of experiences:
        $e_i = (s_i, a_i, s_{i+1}, r_i)$  from replay buffer  $\mathcal{R}$ 
12:    Train critic:
       compute  $y_i = r_i + \gamma Q'(s_{i+1}, \mu'(s_{i+1}|\theta^{\mu'})|\theta^{Q'})$ 
       compute loss  $L = \frac{1}{N} \sum_i (y_i - Q(s_i, a_i|\theta^Q))^2$ 
       perform stochastic gradient descent step on  $L$ 
13:    Train actor:
        $\nabla_{\theta^\mu} J \approx \frac{1}{N} \sum_i \nabla_a Q(s_i, \mu(s_i|\theta^Q)) \nabla_{\theta^\mu} \mu(s_i|\theta^\mu)$ 
14:    update target actor and critic networks parameters:
        $\theta^{Q'} \leftarrow \tau \theta^Q + (1 - \tau) \theta^{Q'}$ 
        $\theta^{\mu'} \leftarrow \tau \theta^\mu + (1 - \tau) \theta^{\mu'}$ 
15:   end for
16: end for

```

Although the interaction of the agent with the environment is similar to DDQN, the exploration in continuous action space is more challenging as the number of possible actions is infinite. The DDPG algorithm uses temporally correlated noise generated using the Ornstein-Uhlenbeck process \mathcal{N} as the exploration technique [Uhlenbeck 1930]. In line 6, the noise value \mathcal{N}_t is added to the action itself, while the process \mathcal{N} is reset at the beginning of every episode. The critical difference of the DDPG algorithm from the DDQN lies in its training process as both actor and critic networks have to be trained. In line 12, the update of the critic network Q is based on the Bellman equation, where Q -values of the next state s_{i+1} are calculated with the target actor network μ' and the target critic network Q' . Afterward, mean-squared loss L is minimized using Adam optimizer with a learning rate α_c . Therefore, the critic training process is similar to the training process of the DDQN algorithm. The actor training, which is done successively, aims at maximizing the Q -values predicted by the critic based on the actions suggested by the actor itself. According to line 13, the gradients of the critic network Q output computed with respect to input actions are fed back to the actor

network μ using the chain-rule to update the parameters of the actor network μ . Once the training is performed, the weights of the target networks $\theta^{Q'}$ and $\theta^{\mu'}$ are slowly updated based on the weights of the main networks, as seen in line 14, according to the averaging rate τ .

Table 9.3 presents the original and empirically chosen hyperparameters of the DDPG algorithm. The Ornstein-Uhlenbeck process used for exploration is deployed with the following parameters: $\sigma = 0.2, \mu = 0, \theta = 0.15$, which are identical to those in the original paper [Lillicrap 2015]. The networks' configuration is 32-32 for both critic and actor networks. One has to note that the learning rate of critic α_c is always bigger than the learning rate of the actor α_a . Moreover, a previous study [Liessner 2019] on the RL for EV management has indicated that the DDPG algorithm is susceptible to the choice of hyperparameters and frequently, even a single wrongly chosen parameter can sabotage the learning process.

Table 9.3 – Hyperparameters of the DDPG algorithm

Hyperparameter	Symbol	Original value	Our value
Replay buffer size	\mathcal{R}	1×10^6	5×10^5
Minibatch size	N	64	64
Discount factor	γ	0.99	0.99
Averaging rate	τ	0.001	0.01
Learning rate actor	α_a	10^{-4}	10^{-4}
Learning rate critic	α_c	10^{-3}	10^{-3}

P-DQN for parametrized action space

Parametrized action spaces are usually either discretized into approximate finite sets or relaxed into continuous sets. However, several attempts were made to develop RL algorithms specifically for hybrid action spaces. These works include Q-PAMDP algorithm presented in [Masson 2015], hybrid actor-critic architecture demonstrated in [Fan 2019], extended DDPG shown in [Hausknecht 2015], and the P-DQN algorithm suggested in [Xiong 2018]. In this thesis, the latter algorithm is chosen for solving the parametrized EV charging control problem due to its intuitive implementation, as it encompasses the properties of both DQN and DDPG techniques and off-policy nature.

Algorithm 3 describes the P-DQN implementation according to [Bester 2019], where the original methodology was enhanced with target networks to improve stability during training. The action a_t in parametrized action space A is defined as a tuple (k_t, x_{k_t}) . The prior $k_t \in K$ represents the discrete action from the set of all discrete actions K , while the latter x_{k_t} represents the continuous parameters associated with this specific discrete action. The P-DQN algorithm is based on two main networks: action value network Q and action parameter network μ . These networks, together with their target copies and a replay buffer \mathcal{R} , are initialized in line 1. The μ network takes the current state s_t as input and produces the parameters x_k for all actions $k \in K$. The Q network outputs the Q-values for all actions k while

taking the state s_t and the parameters x_k as an input. Afterward, in line 8 the desired action a_t is determined by the ϵ -greedy policy. Notably, the P-DQN algorithm deploys two different exploration techniques. In addition to the ϵ -greedy that helps explore the discrete part of the action space, the noise process \mathcal{N} , namely the Ornstein-Uhlenbeck, is used to explore the continuous part. Lines 14 and 15 show the training process of the action value Q and the action parameter μ networks, which is similar to the training process of the DDPG. Once the training is complete, the weights are slowly copied from the main networks to target networks Q' and μ' , as depicted in line 16.

Algorithm 3 Parametrized Deep Q-Networks Learning (P-DQN)

```

1: Initialize: action value network  $Q(s, x_k | \theta^Q)$  and action parameter network  $\mu(s | \theta^\mu)$ , target action value network  $Q'$  and target action parameter network  $\mu'$  with weights  $\theta^{Q'} \leftarrow \theta^Q, \theta^{\mu'} \leftarrow \theta^\mu$ , replay buffer  $\mathcal{R}$ 
2: for each episode do
3:   observe current state  $s_t$ 
4:   initialize random process  $\mathcal{N}$  for action parameter exploration
5:   for each step in the environment do
6:     compute action parameters  $x_k \leftarrow \mu(s_t | \theta^\mu) + \mathcal{N}_k$ 
7:     compute action values  $Q_k \leftarrow Q(s_t, x_k | \theta^Q)$ 
8:     select action  $a_t = (k_t, x_{k_t})$  according to the  $\epsilon$ -greedy policy
9:     execute action  $a_t$ 
10:    observe next state  $s_{t+1}$  and reward  $r_t = R(s_t, a_t)$ 
11:    store  $(s_t, a_t, s_{t+1}, r_t)$  in replay buffer  $\mathcal{R}$ 
12:    update current state  $s_t \leftarrow s_{t+1}$ 
13:    sample minibatch  $N$  of experiences:
        $e_i = (s_i, a_i, s_{i+1}, r_i)$  from replay buffer  $\mathcal{R}$ , decompose  $a_i$  into  $k_i$  and  $x_{k_i}$ 
14:    Train  $Q$ :
       compute  $y_i = r_i + \gamma \max_{k_i \in [K]} Q'(s_{i+1}, \mu'(s_{i+1} | \theta^{\mu'})) | \theta^{Q'}$ 
       compute loss  $L = \frac{1}{N} \sum_i (y_i - Q(s_i, x_{k_i} | \theta^Q))^2$ 
       perform stochastic gradient descent step on  $L$ 
15:    Train  $\mu$ :
       compute loss  $L = -\frac{1}{N} \sum_i Q(s_i, \mu(s_i | \theta^\mu) | \theta^Q)$ 
       perform stochastic gradient descent step on  $L$ 
16:    update target action value and target action parameter networks parameters:
        $\theta^{Q'} \leftarrow \tau \theta^Q + (1 - \tau) \theta^{Q'}$ 
        $\theta^{\mu'} \leftarrow \tau \theta^\mu + (1 - \tau) \theta^{\mu'}$ 
17:   end for
18: end for

```

The hyperparameters are collected in Table 9.4. The network configurations are 32-32 for both action value Q and action parameter μ networks. The parameters of the Ornstein-Uhlenbeck exploration process are identical to those of the DDPG algorithm. The ϵ -greedy exploration process starts with ϵ value of 0.5 and decays exponentially until 0.01.

Table 9.4 – Hyperparameters of the P-DQN algorithm

Hyperparameter	Symbol	Original value	Our value
Replay buffer size	\mathcal{R}	1×10^6	5×10^5
Minibatch size	N	64	32
Discount factor	γ	0.99	0.99
Averaging rate	τ	0.01	0.001
Learning rate action value	α_a	10^{-4}	10^{-3}
Learning rate action param	α_{ap}	10^{-5}	10^{-4}

Hindsight experience replay (HER)

To enable sample-efficient learning from sparse and binary rewards, the HER technique developed in [Andrychowicz 2017] is applied. On the one hand, unlike humans, the model-free off-policy algorithms presented have difficulties learning from failures. On the other hand, the sparse reward structure makes successes rare, thus complicating the training process. The HER technique aims to extract useful information from past environment transactions using the notion of goals. At the beginning of each episode, the desired goal g_d is sampled along with the observation of the current state s_t . In the current case, g_d can be the requested SOC at the EV departure, the desired level of PV self-consumption, or both if the g_d is formulated in the form of a tuple. The whole episode $s_{t_0}, s_{t_1}, \dots, s_T$, where s_T is the state at the episode's terminal, is stored in the form of experiences s_t, a_t, s_{t+1}, r_t in the transaction buffer \mathcal{D} . At the end of the episode, the agent arrives to some achieved goal g_a , which can be identical to g_d , if the episode was a success, or can be different if it was a failure. Here the HER technique comes into play to efficiently utilize the experiences of the episode even if g_d was not achieved. The information stored in the \mathcal{D} buffer is used to replay the whole episode pretending that the achieved goal g_a was the desired goal from the beginning. Therefore, the failure episode becomes a success episode, where the agent acquires the knowledge, learns to achieve the goal g_a , and obtains the reward for successful completion of the task. The newly replayed episode is stored in the replay buffer \mathcal{R} , where it serves as a positive example to train the agent. In the current implementation of the HER technique that follows the original paper [Andrychowicz 2017], the strategy *final* was used, which compares the achieved and desired goals at the terminal state of the episode. The addition of HER technique requires the specific implementation of the environment where the goal is concatenated to the state.

Expert demonstrations

The inclusion of expert demonstrations is another technique suited for facilitating the training process by improving the exploration of the environments with sparse rewards. In this work, expert demonstrations are used to enhance the learning and speed up the convergence of the DDPG algorithm in particular. The main reasons include the tedious exploration of infinite continuous action space and the difficulty of tuning the DDPG algorithm's hyper-

parameters, previously mentioned in the respective subsection. Moreover, earlier successful implementations of the DDPG algorithm leveraging expert demonstrations in robotics [Nair 2018, Vecerik 2017] and in autonomous driving [Zuo 2017] motivate such choice.

The main idea to combine expert demonstrations with RL is to aid the agent's training by providing experiences where the outcome was successful. The practical implementation of DDPG with expert demonstrations consists of three main steps. First, the second replay buffer \mathcal{R}_D is created to store the demonstration data in addition to the original replay buffer \mathcal{R} that holds the self-generated data. Both data are stored in the form of experience tuples (s_i, a_i, s_{i+1}, r_i) . Second, during the training, the minibatch N of previous transactions is sampled from both replay buffers \mathcal{R}_D and \mathcal{R} according to the predefined proportion. Third, the loss L computed during the actor's training is augmented with the behavior cloning loss L_{BC} according to [Nair 2018]. The additional loss is applied only on demonstration samples and is calculated as follows:

$$L_{BC} = \sum_{i=1}^{N_D} \|(\mu(s_i)|\theta^\mu) - a_i\|^2 \quad (9.13)$$

Where N_D is the amount of demonstration samples in the minibatch N , a_i is the action taken by the expert and $\mu(s_i)|\theta^\mu$ is the action produced by the agent's policy. Therefore, behavior cloning loss L_{BC} is combined with the loss of the actor using λ_1 and λ_2 weights:

$$L = \lambda_1 \nabla_{\theta^\mu} J - \lambda_2 \nabla_{\theta^\mu} L_{BC} \quad (9.14)$$

9.4 Case study

To validate the proposed methods for solving the EV charging control problem, a case study was created within the smart community framework proposed in Chapter 7. The agent is an EV defined by its battery capacity and the upper limit on the charging power P_{max} . Similar to Chapter 8, it is the Citroen C-Zero model with 16 kWh battery and 3700 W maximum power input [EVD 2021a]. In addition, the efficiency of the EV charging process is assumed $\eta^{EV} = 1$.

9.4.1 Datasets

The test dataset was collected from one of the hotels participating in the project. The dataset spans over a period of 2 months from the 9th of September to the 9th of November 2020. The building load demand L_t and the PV production PV_t measurements were down-sampled to hourly resolution. Due to the absence of the real-world recordings of the EV-related part of the dataset, EV usage patterns were simulated. The SOC at arrival, time of arrival at the charging station, and time of departure from the charging station were sampled from uniform distributions in the following way:

- $SOC(t = t_{arr}) \sim \mathcal{U}(0.2, 0.5)$

- $t_{dep} \sim \mathcal{U}(7, 19)$
- $t_{arr} \sim \mathcal{U}(t_{dep} + 1, 23)$

According to chosen distribution boundaries, a minimum trip duration of one hour and EV return to the hotel on the day of departure are assumed. The time to EV departure from charging station D_t , which is necessary for defining the state s_t , was computed as the number of hours between the EV arrival at t_{arr} and the next departure t_{dep} .

Due to the limited availability of data from the participating hotel that serves as the smart building, the training dataset for RLC algorithms was synthesized based on another hotel in the area. The smart meter measurements were scaled to reflect the test hotel's size, while the PV production of a virtual installation was simulated using relevant recordings from a nearby meteorological station. The EV driving scenario was generated using the same uniform distributions as shown above. The length of a training dataset is one year.

9.4.2 Performance evaluation

The performance of DDQN, DDPG, and P-DQN algorithms is evaluated in two phases. First, the results obtained during training are presented. Following the established goals, both the PV self-consumption (Equation 9.3) and SOC at departure (Equation 9.4) metrics are presented. To enhance comprehension of the results, the optimal PV self-consumption is computed, according to Equation 9.9. The addition of this metric gives better judgment on the PV usage by the EV charging process. To ensure fairness of reporting deep RL algorithms' performance, the guidelines indicated in [Henderson 2017] are used. Thus, five experiments are run for each of the algorithms using the same set of random seeds. Further, the DDQN, DDPG, and P-DQN algorithms' performance is summarized by reporting the evolution of mean and standard deviation throughout the training process for each of the objectives across five seeds. Moreover, to provide judgment on the complexity of suggested deep RL algorithms and the computational resources required, the duration of the training process is reported.

In the second phase, both benchmark and deep RLC algorithms are executed on the held-out test set of historical data consisting of 60 episodes. For all algorithms the performance metrics, namely PV self-consumption and SOC at departure, are reported in the form of a box-plot visualization that provides a five-number summary: the minimum, the maximum, the sample median, and the first and the third quartiles across 60 episodes. Moreover, the amount of energy purchased from the grid is reported in the same format along with its distribution. To highlight the differences in the EV charging approaches suggested by the algorithms, the evolution of the SOC across episodes with the same SOC at arrival and across episodes with binned duration is compared. Lastly, the information on the execution time of the algorithms on the test set is reported to understand the suitability for online implementations.

9.4.3 Benchmark algorithms

To understand the quality of the solutions provided by RLC methods, they are compared among themselves and with other classic approaches used in the EV charging control problem. In this section, several algorithms used to benchmark the proposed methodology are presented. Importantly, all benchmark solutions have to respect the SOC (Equation 9.5) and the power limit (Equation 9.6) constraints introduced in Section 9.3.1.

Naive EV charging

The naive algorithm reflects the conventional charging attitude of the majority of EV drivers nowadays. The EV starts to charge as soon as it is plugged to the charging station at t_{arr} . If there is any PV surplus, the EV draws the available PV power EV_t^{PV} , otherwise the power supply EV_t^G comes from the grid and is equal to P_{max} .

RBC EV charging

Whenever there is an excess of PV production, the EV uses the renewable power to charge. If the PV is not available, the algorithm evaluates the possibility of drawing power from the grid based on the EV's SOC_t and the time left before departure from the charging station D_t . Based on the following expression, the EV compares the leftover energy to be charged in the battery with the maximum possible transmittable energy during the remaining time.

$$(SOC_{max} - SOC_t)C_{bat} \leq (D_t - t_{lag})P_{max} \quad (9.15)$$

The time parameter t_{lag} tunes how shortsighted the algorithm is. A large value of t_{lag} helps the control to anticipate the EV's departure, however, it should not exceed the time required to fully charge the EV from SOC_{min} to SOC_{max} using P_{max} from the grid. As a result, the EV draws P_{max} from the grid at time t if the EV departs rather soon and chooses not to charge otherwise.

Deterministic optimization of EV charging

The deterministic optimization of EV charging, aimed to provide the optimal solution, is formulated with the following objective function that encompasses both, maximization of the PV self-consumption and SOC at departure, where k_1 and k_2 are respective weights for the two objectives, such as $k_1 + k_2 = 1$:

$$\max \left(k_1 \left(\frac{\sum_{t=t_s}^{t_e} (EV_t^{PV} + L_t^{PV})}{\sum_{t=t_s}^{t_e} PV_t} \right) - k_2 (SOC_{max} - SOC_{t_{dep}}) \right) \quad (9.16)$$

subject to constraints for $t = t_s, \dots, t_e$:

- ▷ $PV_t + y_t^G P_t^B - (1 - y_t^G) P_t^S - EV_t - L_t = 0$,
where P_t^B is the power bought from the grid, P_t^S is the power sold back to the grid, and $y_t^G \in \{0, 1\}$ is the binary variable that forbids the simultaneous purchasing and selling of the power from and to the grid
- ▷ $PV_t = EV_t^{PV} + L_t^{PV} + P_t^S$
- ▷ $P_t^B = EV_t^G + L_t^G$
- ▷ $EV_t = y_t^{EV} EV_t^G + (1 - y_t^{EV}) EV_t^{PV}$,
where $y_t^{EV} \in \{0, 1\}$ is the binary variable that governs the simultaneity of power supply to EV according to Equation 9.2, such as $y_t^{EV} \leq y_t^G$
- ▷ $SOC_{t_{arr}} = SOC_{arr}$
- ▷ Load satisfaction constraint (Equation 9.1)
- ▷ SOC limits (Equation 9.5)
- ▷ EV charging power limit (Equation 9.6)
- ▷ SOC continuity (Equation 9.7)

Notably, the deterministic optimization of EV charging and presented further deterministic and stochastic MPC formulations were developed in a master thesis conducted at PV-Lab. Thus, one can find more detailed explanations of these benchmarking algorithms in [Martinson 2021].

Deterministic MPC of EV charging

The deterministic formulation of MPC is iterative, thus the algorithm takes charging decisions at every time step t . The episode t_s to t_e is separated into two successive stages: the decision stage, where the actual time step $t = t_d$, and the prediction stage $t_f = t_{d+1}$ to t_e . At the decision stage the PV production PV_{t_d} and the load consumption L_{t_d} values are known, while at the prediction stage the Long Short-Term Memory (LSTM) neural network and the Auto Regressive Integrated Moving Average (ARIMA) model are used to predict the PV_T and L_T values for $T = t_f, \dots, t_e$ respectively. As discussed in [Martinson 2021], the statistical ARIMA model that relies solely on historical data is chosen for L_T forecasting due to the lack of any exogenous variables. Instead, the PV_T prediction benefits from the presence of additional meteorological data, such as temperature, radiation, and pressure. Thus, LSTM model, previously successfully deployed in Chapter 5, is used to forecast PV production time-series due to its innate ability for feature selection. The objective function, though similar to Equation 9.16, is transformed in the following way to reflect the two-stage formulation:

$$\max \left(k_1 \left(\frac{EV_{t_d}^{PV} + L_{t_d}^{PV}}{PV_{t_d}} + \frac{\sum_{t=t_f}^{t_e} (EV_t^{PV} + L_t^{PV})}{\sum_{t=t_f}^{t_e} PV_t} \right) - k_2 (SOC_{max} - SOC_{t_{dep}}) \right) \quad (9.17)$$

The constraints applied to the model are identical to those for deterministic optimization stated previously and are formulated for both decision and prediction stages.

Stochastic MPC EV charging

The stochastic formulation of MPC differs from the deterministic formulation during the prediction stage, as LSTM and ARIMA models forecast confidence intervals for PV_T and L_T . Therefore, one can draw an Ω set of N scenarios from predicted distribution, where each scenario $w_i \in \Omega$ consists of $PV_T^{w_i}$ and $L_T^{w_i}$ forecast values. All constraints stated for deterministic MPC are equally true for the stochastic MPC and require two-stage formulation. To take into account all possible scenarios from Ω set, the objective function has to be reformulated in the following way, assuming that each $w_i \in \Omega$ scenario is equally probable:

$$\max \left(k_1 \left(\frac{EV_{t_d}^{PV} + L_{t_d}^{PV}}{PV_{t_d}} + \frac{1}{N} \sum_{i=0}^N \frac{\sum_{t=t_f}^{t_e} (EV_t^{PV, w_i} + L_t^{PV, w_i})}{\sum_{t=t_f}^{t_e} PV_t^{w_i}} \right) - k_2 \frac{1}{N} \sum_{i=0}^N (SOC_{max} - SOC_{t_{dep}}^{w_i}) \right) \quad (9.18)$$

9.4.4 Implementation

To apply RLC for the EV charging control problem, both the environments and the algorithms were implemented using Python programming language. Custom discrete, continuous, and parametrized environments were created with the OpenAI Gym toolkit [Brockman 2016]. A specific *gym.GoalEnv* class was chosen to build goal-based environments to enable the HER technique. The RL algorithms were developed using *tensorflow 2.0* library. The deterministic optimization and both stochastic and deterministic MPC algorithms rely on *Gurobi* solver. All the experiments were conducted using a personal laptop (Intel i7- 7600, 16 GB RAM).

9.5 Results

9.5.1 Training phase

Figure 9.3 depicts DDQN, DDPG, and P-DQN algorithms' performance during the training phase of 2000 episodes on a one-year dataset. All algorithms deploy the HER technique to facilitate the training, while the DDPG algorithm's learning process is additionally enhanced by 100 episodes of expert demonstrations. The bold line and the shaded region on each subplot show the mean and the standard deviation of 5 runs, each corresponding to a different seed. The evaluation of algorithms with various seeds is important to understand the models' robustness, validate the choice of the hyperparameters, and compare the algorithms among themselves. Although seeding can fix the random processes to get reproducible results in controlled experiments, one has to note that randomness is an inherent feature of ML models.




Figure 9.3 – Training performance of deep RL algorithms

Therefore, when commissioning models in production, one should remove the fixed seed value. During the training process, the following major parts of the RL workflow that introduce randomness were identified:

- Arrival SOC initialization at the beginning of each episode
- Weight initialization in neural networks
- Adam optimizer
- Ornstein-Uhlenbeck and ϵ -greedy exploration policies
- Minibatch sampling from replay buffer \mathcal{R}

As shown in Figure 9.3, the P-DQN algorithm converges to higher optimal PV self-consumption, followed by DDPG and DDQN, with a substantial share of episodes reaching the maximum PV self-consumption possible. The same metric for DDPG stabilizes to 0.9 with a very low standard deviation, which can be explained by the inclusion of behavior cloning loss from expert demonstrations along the training process. The DDQN algorithm shows interesting behavior, as PV self-consumption significantly changes its variance and starts to fluctuate between 0.8 and 0.95 from approximately 1100 episodes. The turning point corresponds to the ϵ -value of 0.3 in the ϵ -greedy policy, which decays exponentially throughout the training process. The same behavior of the DDQN algorithm can be noted for the SOC at departure metric. The standard deviation across runs increases, while the mean SOC at departure decays to an average of 85%. However, during the whole training process, a substantial share of the episodes achieves SOC_{max} at departure. The P-DQN algorithm exhibits high variance as well, with the mean SOC at departure at 80%. The DDPG algorithm converges to the highest

SOC at departure of 95% among all algorithms with very low variance. Moreover, the DDPG demonstrates the fastest speed of learning due to facilitation through expert demonstrations. Particularly, the experiments show that the inclusion of expert data improves DDPG's PV self-consumption and SOC at departure by 0.02 and 0.05, respectively, compared to DDPG implementation with HER. As for the influence of HER technique, adding it to algorithms' training processes without any enhancements increases PV self-consumption of DDQN and P-DQN by 0.06 on average and shows no significant effects in case of DDPG. Regarding SOC at departure, implementation of HER has no influence on DDQN, while P-DQN algorithm benefits from 0.03 improvement. Similarly, DDPG demonstrates an increase of SOC at departure by 0.06 compared to the implementation without any special facilitation techniques. Thus, one can conclude that the inclusion of HER leads to marginal improvements in performance for almost all algorithms, while the addition of expert demonstrations to DDPG boosts PV self-consumption and SOC at departure even further. Moreover, the observation of training runs has shown that implementing HER prevents sudden degradation of DDPG and P-DQN algorithms' performance, referred to as catastrophic forgetting related to neural networks. The training results overall confirm that deep RLC algorithms are capable of achieving multiple objectives simultaneously. The maximization of the PV self-consumption is more pronounced for DDQN and P-DQN algorithms, although the SOC at departure values show a decaying trend. However, the mean values of SOC at departure for these algorithms are exceeding the 80% limit, which is considered acceptable SOC before departing for the majority of EV drivers.

Figure 9.4 depicts the training duration of deep RLC algorithms. The DDQN algorithm is shown twice as the averaging rate τ and the update every n hyperparameters influence the training duration differently. With $\tau = 0.01$, the online neural networks are trained at every step $n = 1$, and the weights are slowly copied to the target networks. With $\tau = 1$, the training is performed every $n = 32$ steps; thus, the overall process takes considerably less time.




Figure 9.4 – Training duration of RL algorithms

There are three key observations that one can make from Figure 9.4. First, the training duration is the linear function of the number of episodes, which is useful to estimate required time resources before engaging in training. Second, the execution time varies for the same number of episodes for the same algorithm as each episode's duration is different, and the environment is reset at the beginning of each episode. Third, the DDPG algorithm requires almost twice longer training time than other algorithms due to the complexity of its neural networks framework, larger minibatch size, and additional computation of behavior cloning loss.

9.5.2 Testing phase

During the testing phase, the deep RLC algorithms are compared with benchmark algorithms chosen in Section 9.4.3 on a held-out test dataset of 60 episodes. The current section presents the results in ascending order of time granularity, starting with observations per episode and following with comparison across the whole test dataset.

Figure 9.5 depicts the EV charging strategies for all algorithms on three sample episodes. The episodes were chosen to demonstrate the diversity of PV production and load consumption profiles. The deterministic algorithm demonstrates the optimal performance in terms of both, the amount of PV consumed for EV charging and the $SOC_{t_{dep}}$. Despite the Naive algorithm always supplying the grid power for EV charging, as the EV's arrival times are outside the PV production hours, the algorithm guarantees the $SOC_{t_{dep}} = SOC_{max}$. The RBC algorithm exhibits the charging strategy similar to the optimal performance, though some of the grid power inputs are shifted in time. Both deterministic and stochastic MPC algorithms show good usage of PV generation and achieve rather high $SOC_{t_{dep}}$ for all three episodes. However, the sudden load consumption peaks, such as at the end of the episode 1, are not handled well to provide SOC_{max} at departure. The deep RLC algorithms demonstrate tendency towards using the most PV power possible and in two out of three episodes reach high $SOC_{t_{dep}}$. However, the absence of the PV power production period in the episode 3 prevents all RLC algorithms from ending the episode with high SOC, while the short duration of the episode leaves less time for decision-making. The DDPG's grid power supply strategy differs from all other algorithms. Particularly, the DDPG takes lower risks by purchasing electricity early and in small quantities, thus anticipating potential future disturbances in PV power production.

The boxplots in Figure 9.6 summarize the algorithms' performance across all episodes in the test dataset. The deterministic algorithm, indeed, serves as the baseline for comparison as it provides the optimal performance. The Naive algorithm demonstrates the lowest PV self-consumption values among algorithms while it guarantees the $SOC_{t_{dep}} = SOC_{max}$. The RBC algorithm performs closely to the optimal solution, though with almost twice higher variance in PV self-consumption than the deterministic algorithm. The stochastic MPC outperforms the deterministic MPC in PV self-consumption, while in some cases it does not reach the SOC_{max} at departure. The deep RLC algorithms vary significantly in their $SOC_{t_{dep}}$ levels from all other algorithms, with the DDPG having the highest $SOC_{t_{dep}}$ values among RLC

algorithms. The DDQN algorithm demonstrates mediocre performance for both objectives. The P-DQN algorithm achieves a higher median for PV self-consumption than for $SOC_{t_{dep}}$, though is characterized by very high variance across episodes. The median values of energy purchased from the grid are very similar among all algorithms, with DDQN and DDPG using less grid energy. However, marginal increases of the purchased energy amounts for these algorithms would possibly compensate for lower $SOC_{t_{dep}}$. Thus, one can conclude that major discrepancies between EV charging strategies result primarily from differences in PV usage.

Figure 9.7 shows density histogram plots of grid power usage for EV charging among the algorithms. Although the flattening of the grid-supplied part of the EV charging curve was not specified as one of the control objectives, the difference in grid power utilization can be observed among algorithms. Particularly, one can notice that both Naive and RBC algorithms rely heavily on P_{max} injections to the EV. Instead, the power of smaller magnitudes is used occasionally and is more probable to appear closer to the end of the charging process as the SOC increases. The deterministic and the stochastic MPC have similar grid power usage patterns, with the prior supplying higher power more frequently. On the contrary, the deterministic MPC has a more balanced profile of EV supply from the grid, with minimal and substantial power values being equally probable. Although the DDQN's grid power usage pattern is similar to that of non-RLC algorithms, the DDPG and P-DQN demonstrate different approaches to drawing power from the grid for EV charging. Particularly, the DDPG algorithm relies only on small and frequent grid power quantities and does not induce any high power peaks during the charging process. Such behavior can be explained by the DDPG's anticipation and risk minimization strategy. Similarly, the P-DQN algorithm uses small-magnitude grid power values while rarely injecting P_{max} to the EV.

Figure 9.8 depicts the SOC evolution across episodes with $SOC_{t_{arr}} \in \{20\%, 30\%, 40\%, 50\%\}$, where the speed of SOC changing and the differences between the charging approaches undertaken by considered algorithms can be observed. The lengths of the episodes were normalized to provide the possibility to compare the SOC profiles across charging processes with various duration. The Naive algorithm shows a steep profile of SOC evolution reaching SOC_{max} within initial 20% of the charging time regardless of the SOC_{arr} . Instead, the RBC algorithm starts to charge the EV closer to the end of the charging process. The deterministic approach and both MPC algorithms have less steep SOC evolution profiles, with the charging process starting later for EVs arriving with $SOC_{arr} = 50\%$ than for EVs with lower SOC_{arr} . However, one can notice that with $SOC_{arr} = 30\%$ almost the whole duration of the episode is used to charge the EV. The DDPG algorithm stands out with its linear SOC profile and effective usage of the whole episode for charging. Moreover, one can notice that the cases of not achieving the SOC_{max} at departure are more common for $SOC_{arr} = 20\%$. The P-DQN algorithm increases its waiting time before charging according to SOC_{arr} increases, while the same feature does not appear in the DDQN's SOC evolution profile.



Figure 9.5 – Examples of benchmark and RLC algorithms on selected episodes of test data




Figure 9.6 – RLC and benchmark algorithms’ comparison across episodes




Figure 9.7 – Comparison of grid power usage across algorithms

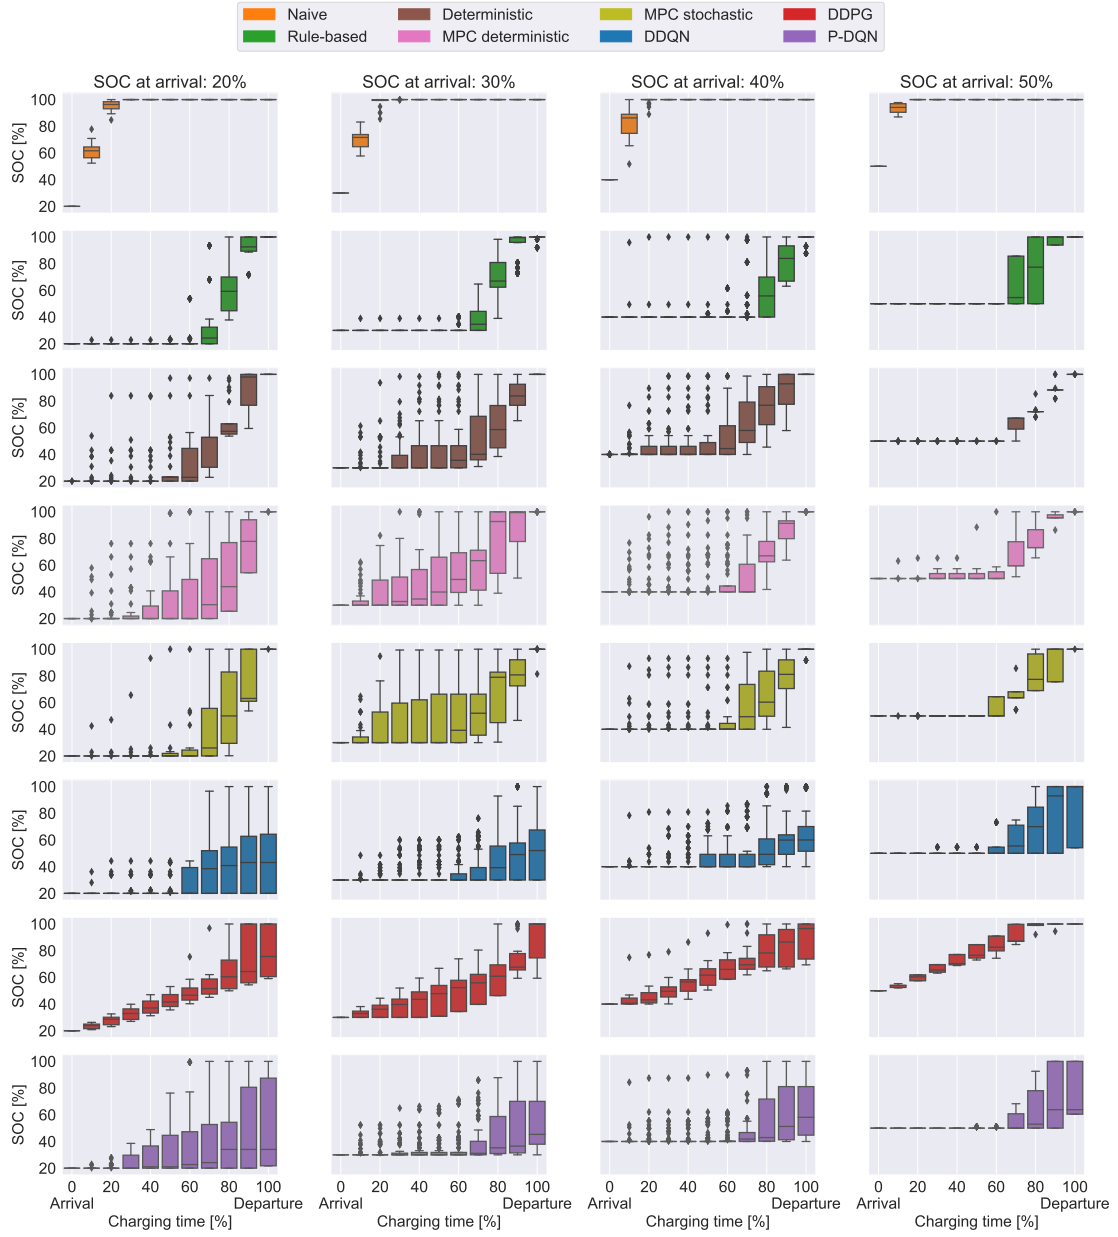


Figure 9.8 – Comparison of SOC evolution across episodes with equal SOC at arrival

Figure 9.9 shows the SOC evolution, PV generation, and load demand profiles across episodes with binned duration. One can notice that longer episodes feature two peaks of PV generation and load consumption. While the EV charging strategies of Naive and RBC algorithms do not exhibit noticeable changes with increased episode's length, the deterministic and both MPC algorithms demonstrate interesting changes in the SOC profile. The variance during the initial 20% of the charging time can be explained by the spread of the SOC_{arr} . The SOC profiles for short episodes below 14 hours exhibit logarithmic growth, while the episodes with 19-26-hour duration show exponential growth. The latter can be explained by weak PV generation profiles during initial 50% of the charging time. The shape of the SOC curve for long episodes changes to accommodate a double-period of PV generation. Thus, the SOC increases initially and at the end of an episode, while the middle part stays relatively steady. A similar double-curvature feature can be noticed for the P-DQN algorithm at very long episodes, while everywhere else, the SOC profile resembles a mixture between linear and exponential. The DDPG algorithm demonstrates a stable linear charging trend regardless of the episode's duration. The DDQN's SOC profile shows efficient usage of PV generation for episodes with 23-26-hour duration. However, for other episodes, the small slope of rather a linear trend results in low SOC_{dep} .

After looking at the algorithms' performance across episodes, the results are summarized on the whole test dataset in the remainder of this section. Table 9.5 expresses the SOC at departure satisfaction levels for each algorithm. The low level of satisfaction corresponds to $SOC_{dep} \leq 50\%$, making the EV unsuitable for usage. The high level of satisfaction conforms with $SOC_{dep} \geq 80\%$ and is achieved by all non-RLC algorithms on 100% of test episodes. The RLC algorithms obtain less SOC_{dep} satisfaction, whereas 15% of DDQN episodes and 21% of P-DQN episodes make the EV unfitted for driving purposes. Instead, the DDPG algorithm does not exhibit any episodes with low SOC satisfaction but its 70% share of high satisfaction episodes is still lower than for non-RLC algorithms. The SOC_{dep} of DDQN and P-DQN are located in the middle and high-satisfaction zones, respectively.

Table 9.5 – Share of episodes with various SOC at departure satisfaction levels

Algorithm	Low [%]	Medium [%]	High [%]
Naive	0	0	100
Rule-based	0	0	100
Deterministic	0	0	100
MPC deterministic	0	0	100
MPC stochastic	0	0	100
DDQN	15	52	33
DDPG	0	30	70
P-DQN	21	32	47

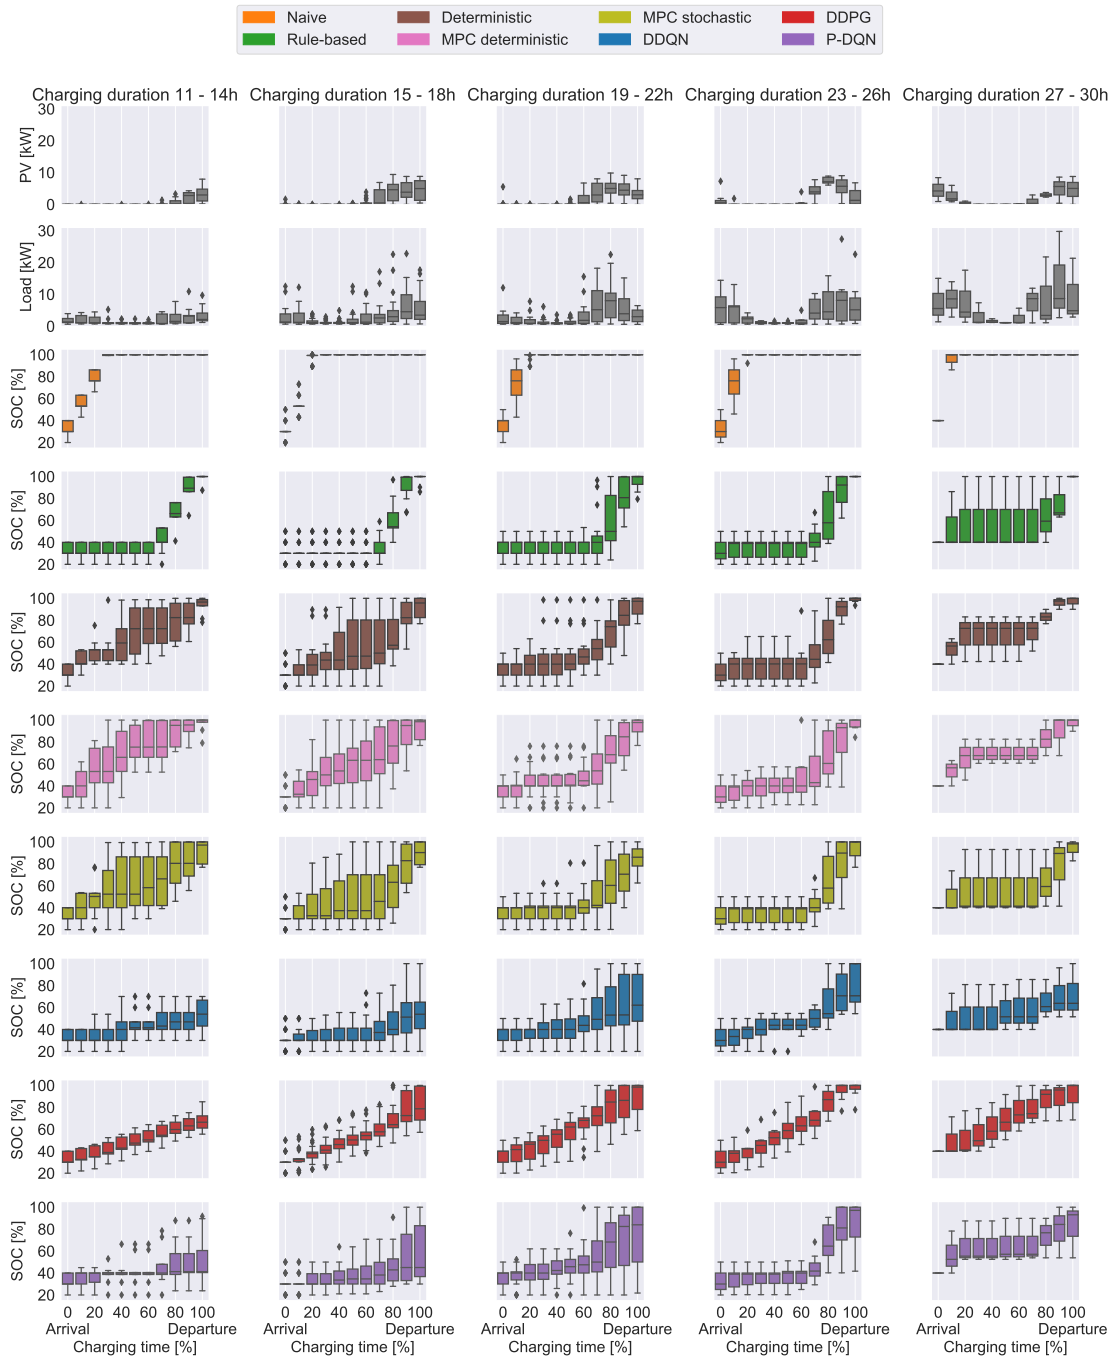


Figure 9.9 – Comparison of SOC evolution across episodes with binned charging duration

Table 9.6 lists the total PV self-consumption and total energy purchased from the grid for all algorithms on the whole test dataset, with a deterministic algorithm setting the baseline for comparison. The Naive algorithm confirms the need for deploying EV charging control by showcasing the lowest PV self-consumption value and the highest amount of energy purchased from the grid. Despite its non-sophisticated nature, the RBC algorithm demonstrates

near-optimal performance, thus justifying its wide-spread usage in control applications nowadays. The MPC stochastic exhibits the closest to the optimal performance by leveraging the power of both PV and load demand forecasts. The MPC deterministic performs in-line with RLC approaches. The P-DQN and DDQN algorithms demonstrate the highest total PV self-consumption among RLC algorithms, while the total purchased energy is the lowest. Therefore, despite the obvious potential for higher PV self-consumption, one can hypothesize that more energy could have been purchased to increase SOC_{dep} . Instead, the DDPG algorithm did not realize its full potential to harvest all available PV generation.

Table 9.6 – Total PV self-consumption and total energy purchased from the grid

Algorithm	PV self-consumption	Energy purchased [MWh]
Naive	0.664	3.359
Rule-based	0.808	3.115
Deterministic	0.829	3.084
MPC deterministic	0.778	3.170
MPC stochastic	0.817	3.100
DDQN	0.762	2.860
DDPG	0.727	3.123
P-DQN	0.774	2.878

Table 9.7 summarizes the execution time of the algorithms on the test dataset. The pre-training duration of the RLC algorithms is not considered as part of the execution time. Once the algorithm is trained offline, it can be efficiently utilized for decision-making. The Naive, RBC, and RLC algorithms execute within seconds, while the MPC deterministic and MPC stochastic require approximately 12 minutes and 1 hour to execute. The latter's execution time makes it unsuitable for online applications, as the chosen time resolution of the problem is equal to 1 hour. The MPC deterministic can be deployed online. However, it requires significantly more computational power than RLC algorithms.

Table 9.7 – The algorithms' execution time on test dataset

Algorithm	Time [s]
Naive	0.21
Rule-based	0.22
Deterministic	21.04
MPC deterministic	772.48
MPC stochastic	3670.66
DDQN	2.16
DDPG	2.27
P-DQN	2.57

To summarize the results described in this section, the key takeaways on the algorithms' performance on the EV charging control problem are highlighted:

- The Naive, conventional charging attitude of the majority of the EV drivers nowadays leads to low PV self-consumption due to extensive charging using the grid-supplied power. Although the latter results in high EV charging costs, the Naive algorithm guarantees SOC_{max} at departure, thus alleviating the range anxiety problem.
- The deterministic optimization algorithm, despite providing the optimal charging strategy, cannot be considered realistic as it assumes complete knowledge of the future PV generation and load demand. Thus, it can be used as a baseline to compare other algorithms' performances but should not be considered a standalone EV charging strategy.
- The RBC algorithm demonstrates near-optimal performance with simple rule formulation, low execution time, and high PV self-consumption and SOC_{dep} values, thus justifying its wide-spread application for control problems nowadays. However, one should carefully consider RBC for problems with increased complexity. The addition of such features as V2G, multiple EVs, and dynamic pricing can result in difficulties in the formulation and verification of the RBC rules.
- The MPC algorithms demonstrate their ability to achieve the PV self-consumption and SOC_{dep} objectives efficiently. However, their formulation requires a full mathematical model of the physical system, development and integration of forecasting instruments, high computational resources, and long execution times to provide decisions. Following the increasing trend of instantaneous decision-making, MPC algorithms' utilization might not respond to the future needs of online implementation. Moreover, the growing complexity of physical systems demands more complicated mathematical models, which increase the execution times even further.
- The RLC algorithms prove their ability to be used for EV charging control despite the varying performance among considered models. With the increasing abundance of data and facilitated access to computing power, one can argue to improve the RLC algorithms' performance in the future significantly. Almost instantaneous decision-making of RLC algorithms leaves a great promise for real-time applications and poses a serious challenge to MPC algorithms with their lengthy execution times. Besides, the increasing complexity of future mobility systems can be efficiently handled by RLC through the utilization of collected big data for learning. Moreover, the model-free nature of RLC method will remain an advantage, eliminating the need for extremely complex and tedious to formulate mathematical models.

9.6 Conclusion

To conclude, this chapter demonstrated the application of RL to the EV charging control problem, focusing on a simple energy system composed of a utility grid, building load, PV

generation, and a single EV. Particularly, the three mathematical formulations of the problem in the form of MDPs that differ by the type of action space were introduced. Moreover, the pool of EV charging control objectives seen in the literature was extended by focusing on maximizing PV self-consumption and EV SOC at departure simultaneously. To resolve the suggested MDP formulations, the DDQN, DDPG, and P-DQN RLC algorithms were deployed for discrete, continuous, and parametrized action spaces, respectively. Throughout a comprehensive benchmarking procedure conducted on a held-out test dataset, the RLC method was compared with naive, RBC, deterministic optimization, and MPC deterministic and stochastic approaches. The comparison has shown that despite a slightly lesser performance of RLC algorithms on the chosen objectives, the RLC approach has exhibited its consistency in delivering applicable EV charging control strategies. Moreover, RLC's execution within seconds has demonstrated a great potential for efficient online implementations, which would be difficult in the case of deterministic and stochastic MPCs with their execution times in the order of minutes and hours. Furthermore, the RL methodology has demonstrated a natural fit to the growing complexity of future energy systems characterized by the abundance of data. Thus, the RLC of EV charging can take its rightful place among other digital energy services and become an indispensable tool for facilitating EVs' integration in the broad energy system. To improve the functionality of the suggested digital energy service, the following enhancements have to be considered:

- For better convenience of the EV charging service, an interface that facilitates bilateral communications between the EV and the charging station has to be introduced. Particularly, the suggested MDP formulation requires the EV driver to submit the SOC and time to departure information at the arrival to the charging station. Although new-generation chargers are capable of extracting the EV's SOC as soon as an EV is plugged for charging, the time to departure variable has to be submitted by the EV driver through a specifically designed interface, such as a mobile application.
- To fully embrace the EVs' mobile storage capabilities and utilize more efficiently the EV parking time, the V2G and vehicle-to-home options have to be included in the problem's formulation. The suggested parametrized implementation of the action space represents the best fit for incorporating these additional functionalities while considering the necessary constraints on simultaneous purchasing and selling of the electricity from and to the grid.
- To maximize the utility of various additional sources of flexibility, such as heat pumps, ESS, and boilers, they should be included in the energy system's representation to reflect the complexity of real-world microgrids. The OpenAI Gym environments developed in the current chapter can serve as a basis for successfully implementing these components, while the proposed algorithms can still be applicable. However, one would need to redefine the notion of an episode for better usage of additional flexibility sources.
- The search for other RL methods that can successfully merge multiple, sometimes conflicting, objectives has to be continued. Particularly, one can analyze various reward

schemes and focus on specific methodology to determine appropriate weights when fusing several objectives in continuous rewards and deciding on their prioritization. Alternatively, the application of Pareto RL for the multi-objective EV charging problem has to be investigated.

- To enable system's scalability, the suggested deep RLC approach has to be extended to multiple EVs and potentially multiple smart buildings and PV installations. Moreover, one can choose to switch from single-agent to multi-agent methodology to explore how RL agents cooperate towards common or competitive goals.

While the current chapter focused on developing control methods for devising intelligent EV charging strategies, the mechanism that ensures reliable accounting of resulting energy flows was not discussed. Although in the case when one person owns both an EV and a charging station settling energy interactions is not required, the situation becomes complicated when multiple parties get involved. Therefore, the following chapter investigates the potential of applying blockchain technology for ensuring the correctness of energy accounting in a safe, reliable, and decentralized manner.

10 Blockchain Energy Exchanges

Profound changes driven by decarbonization, decentralization, and digitalization are disrupting the energy industry, bringing new challenges to its key stakeholders. In the attempt to address the climate change issue, increasing penetration of renewables and mobility electrification augment the complexity of the electric grid, thus calling for new management approaches to govern energy exchanges while ensuring reliable and secure operations. The emerging blockchain technology is regarded as one of the most promising solutions to respond to the matter in a decentralized, efficient, fast, and secure way. Thus, the blockchain-supported energy management applications are introduced.

The main **highlights** and **contributions** of this chapter are:

- An extensive comparison between some of the most popular blockchain platforms, such as AragonOS, Energy Web Chain, Hyperledger Fabric, and Ethereum, provided across a set of comprehensive criteria, such as deployment, maintainability, and scalability.
- The proposal of a blockchain-based charging management framework for Electric Vehicles (EVs), tightly interlinked with physical infrastructure and implemented in a real-world demonstration site. With a specifically designed smart contract governing the charging process, the implemented framework enables secure and reliable accounting of energy exchanges in a network of trustless peers, thus facilitating the EVs' deployment.
- Evaluating the proposed blockchain framework's transaction latency and fees on a real-world case study highlights the blockchain's cost-reduction potential along with key obstacles to be resolved for further proliferation of blockchain in the energy sector.

Related **publications**:

[8] Marina Dorokhova, Jérémie Vianin, Jean-Marie Alder, Christophe Ballif, Nicolas Wyrsh and David Wannier. *A blockchain-supported framework for charging management of electric vehicles*. Submitted, 2021

10.1 Background and motivation

Since decarbonization, decentralization, and digitalization bring new challenges to the energy sector [Di Silvestre 2018a], the ever-complexifying electric grids call for new management approaches and new business models to empower efficient interactions between consumers, prosumers, grid operators, energy providers, and legal authorities. The growing number of independent power producers questions utilities' exclusivity to provide energy supply, thus requiring to revise the established principles of trust among key energy stakeholders. Consequently, disruptive approaches to govern energy exchanges while adhering to specific requirements, such as decentralization, conflict-resolution mechanism, non-repudiation, efficient data management, privacy protection, and scalability [Hassan 2019], are required to facilitate the transition towards new-generation smart energy communities.

Blockchain is an emerging technology that can support such transition of the energy sector into the digital era by providing secure, fast, transparent, decentralized, efficient, and low-cost operational solutions [Miglani 2020]. Applicable to trustless environments, where neither reciprocal trust between participants nor the presence of central authority are required, blockchain technology paves the way for the energy industry to embrace the benefits of a peer-to-peer world. Moreover, it is regarded as a well-suited tool to achieve key objectives of future development, such as cost reduction by optimizing energy processes, improved energy security, and increased sustainability by facilitating the integration of renewables [Andoni 2019]. Therefore, it comes without a surprise that various use cases of blockchain in energy have emerged, with decentralized control of power grids and demand-side management, peer-to-peer energy trading, EV management, carbon emission trading, and green certificates being among the most prominent ones [Bao 2020].

10.1.1 Blockchain fundamentals

This section presents the basics of blockchain technology and provides a brief overview of smart contracts. Blockchain is a decentralized digitally distributed ledger that stores transactions of a trustless network of peers in a secure manner. Participants in the blockchain are represented as nodes that cooperate to maintain the data stored on the blockchain. The transactions are packaged in blocks of a certain size that are chained in a chronological order using advanced cryptography, thus making the blockchain grow as the time progresses. The representation of a blockchain is depicted in Figure 10.1.

A combination of hashing and encryption is used for securing the various elements of the blockchain. The hash function transforms input data into a unique value of a fixed length so that no one can derive the original content from the hash value. Therefore, the additional security level is added, and the size of the incoming data is reduced. The transactions packaged into the block are hashed pairwise using the Merkle tree until a single hash value remains that forms the block's Merkle root hash. The hash of the block itself is generated by hashing together the timestamp, hash of a previous block, Merkle root, and other important information.




Figure 10.1 – Blockchain concept

Therefore, blocks in a blockchain are interlinked as every block points to the previous block's hash value except a genesis block. Such a link structure makes blockchain immutable, as changing a single data point would entail recalculation of the hash value of the respective block. As a result, all subsequent blocks will become invalid, and the modification will be rejected.

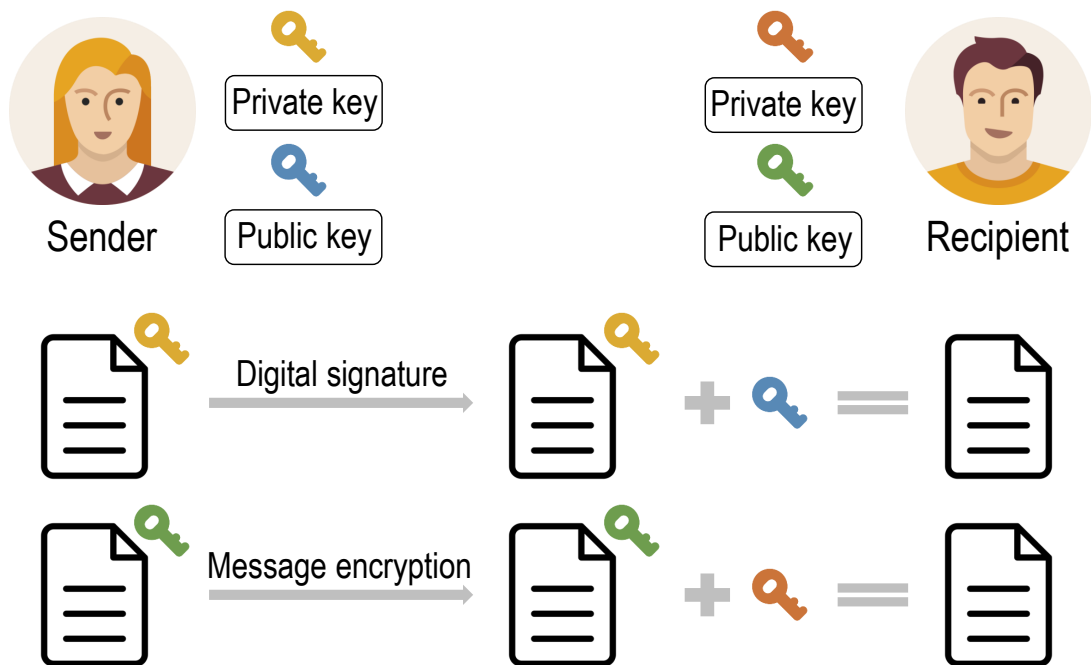


Figure 10.2 – Message encryption and digital signature concepts in blockchain

Blockchain uses asymmetric cryptography to encrypt and decrypt data based on public-private key pairs, as seen in Figure 10.2. The public and private keys are cryptographically linked so that the public key can be derived from the private key using a hash function while the opposite is impossible. The transaction begins with sender and recipient exchanging their public keys, which are essentially their addresses within a blockchain network. Afterward, the sender signs the transaction digitally using a combination of his private key and the transaction's hash value and encrypts the message with the recipient's public key. The latter verifies the sender's identity using a previously received public key and decrypts the message using his own private key.

Depending on the permission model, blockchain technology can be classified into three main types: public, consortium, and private blockchain [Hassan 2019]. A public or permissionless blockchain is an open network accessible to any party, where any node can view, read, and write data on the blockchain. Being the essence of decentralization, public blockchain does not have any association with third parties that act as governing authorities. A consortium or permissioned blockchain is a semi-private network, where a group of organizations controls the ability of a node to join, read, or write to the blockchain. In particular cases, external to consortium nodes can access the content without modification rights to achieve greater transparency. A private blockchain is a restricted type of a permissioned blockchain, where a single organization has full control over the network. Therefore, it offers only partial decentralization as a single entity determines participating nodes and governs the consensus process in the network of trusted parties.

The blockchain node is a critical element of the infrastructure classified as either a normal or a miner node. Although all nodes are responsible for safeguarding the integrity and reliability of the blockchain, their roles vary. Normal nodes can be further divided into full and light nodes. Full nodes are data-heavy as each of them stores a copy of an entire blockchain, including transactions, timestamps, and all created blocks. Their main role is to verify, authenticate, and store the blocks in the network. Light nodes do not hold copies of blockchain as they store only headers of each block, thus depending on full nodes to access complete validated information. Mining nodes are responsible for generating new blocks, broadcasting them to other nodes in the network, and adding them to the blockchain on top of pre-existing nodes once the validation from full nodes is received. As the blockchain technology is decentralized, a consensus mechanism is required to achieve the agreement between nodes on the current state of the network. The consensus protocol's main requirements are being fault-tolerant, energy-efficient, secure, synchronized, low latency, deterministic, resilient to faulty nodes and message delays, and resistant to sophisticated hardware [Miglani 2020]. The most popular consensus mechanisms for blockchain are Proof-of-Work (PoW), Proof-of-Stake (PoS), Proof-of-Authority (PoA), Practical Byzantine Fault Tolerance (PBFT), proof-of-capacity, PoB, and proof-of-luck. In PoW, miners compete in solving complex mathematical puzzles to claim the right to add the block to the blockchain and obtain rewards. Although this process is computationally heavy and energy-intensive, it gives robustness to PoW against malicious attacks as it is almost impossible to acquire 51% of computing power to gain control of

the network. The PoS and PoB are examples of more environmentally-friendly consensus mechanisms that do not require energy-hungry calculations. The prior algorithm favors miners with larger amounts of cryptocurrency at stake, while the latter demands miners to invest coins to an unspendable address.

The advent of Bitcoin blockchain [Nakamoto 2008] in 2009 has transformed the concept of digital payments by removing the middle-man and introducing cryptocurrency. Since then, the scope of blockchain expanded far beyond financial applications when Ethereum decentralized platform [Buterin 2014] established smart contracts previously proposed in [Szabo 1997]. The user-created smart contract is essentially an executable piece of code that resides on the blockchain and runs automatically when predefined initial conditions are met. To make the smart contract functional, one has to define the parties that enter into the agreement, the agreement's subject, and specific terms of the contract. The latter include requirements expected from all participants, mathematical rules that define the contract's enforcement, and rewards and penalties associated with the smart contract's successful run. To execute the smart contract, the operations comprising it are performed by miners within the Ethereum Virtual Machine. Each operation, such as addition or subtraction, is associated with the specific gas cost that measures the computational effort required to carry it. Thus, the total amount of gas needed by the contract depends largely on its size and complexity. The resulting transaction fee can be calculated by multiplying the gas by its unit price, which is determined based on the trade-off between the supply and demand of computing power.

The main advantages of smart contracts are their speed and accuracy, trust, as all nodes in the blockchain witness the execution of a smart contract, savings due to third-party elimination, and security due to the cryptographic basics of blockchain. The disadvantages are uncertain legal status and high implementation costs, as an experienced programmer with expertise in the field is required to establish the agreement's rules correctly. Moreover, the deployment of a smart contract on the blockchain is irreversible, meaning that it cannot be altered once created. Although immutability is the main asset of smart contracts, encoded errors and loopholes can threaten security and lead to hacker attacks as it happened with Decentralized Autonomous Organization (DAO) [Leising 2017], which resulted in a radical change of Ethereum blockchain protocol referred to as a hard fork.

10.1.2 Blockchain glossary

This section compiles the essential terminology in the blockchain field to facilitate the comprehension of the smart contract introduced in Section 10.3 and the blockchain's viability analysis presented in Section 10.6. Thus, the main blockchain-related concepts are defined as follows:

- **Address** is an alphanumeric character string used to send and receive transactions on a blockchain network.
- **Block** is a group of transactions entered into a blockchain.

- **Consensus mechanism** is a method of authenticating and validating a value or transaction on a blockchain network without the need to trust or rely on a central authority.
- **DApp** is a decentralized computer application that exists and runs on a distributed computing system, such as blockchain.
- **Finality** is the state of the blockchain, under which the transaction cannot be canceled, reversed or changed by any of the network's participants under any circumstances.
- **Fork** is a creation of an ongoing alternative version of the blockchain. Soft forks render two chains with some compatibility while hard forks result in a new version of the chain requiring all users to upgrade their clients.
- **Gas** is a measure of the computational steps required for a transaction on the Ethereum network that equates to a fee for network users paid in small ETH units called Gwei.
- **Mainnet** is the primary network where actual transactions take place on a specific distributed ledger.
- **Mining** is the process by which blocks or transactions are verified and added to a blockchain.
- **Node** is a computer which holds a copy of the blockchain ledger.
- **Sharding** is a concept of partitioning a distributed database to allow greater scalability.
- **Smart contract** is a self-executing computer program or a transaction protocol intended to control or document legally relevant events and actions according to the terms of a contract or an agreement.
- **Testnet** is a staging blockchain environment for testing applications before putting them into production or onto the mainnet.
- **Transaction fee** is a small fee charged from blockchain users to process their transaction on the network.
- **Transaction latency** or average transaction time is the time elapsed between the transaction's generation and its final appearance in the block on the blockchain.
- **Transaction Per Second (TPS)** is a measurement of the speed of a blockchain.

The remainder of this chapter is structured as follows. Section 10.2 reviews the key literature describing blockchain's application to energy use cases. Section 10.3 presents the research approach, particularly elaborating on the choice of blockchain and smart contract design. Section 10.4 discusses the system's implementation details and introduces the web interface and mobile application. Section 10.5 presents the real-world case study used to validate the approach and summarizes the main results. Section 10.6 elaborates on the performance of a utilized blockchain. Section 10.7 concludes the proposed blockchain digital service for energy exchanges and suggests further improvements to facilitate its adoption.

10.2 State of the art

The use of blockchain in smart energy systems is a topic of growing research interest over the past few years. The unique features of blockchain, together with specific challenges of the energy sector, motivate researchers to explore new implementation scenarios, define various objectives, and examine diverse energy systems. Previous works, particularly in the EV management application, can be grouped into three categories according to the type of blockchain deployed: public, consortium, and private blockchain.

The first group of studies uses public blockchain, mainly Ethereum, as a transaction mediator in various energy management tasks within smart energy communities. Researchers in [Pustišek 2016] developed an autonomous charging station selection process using the smart contract capability of the Ethereum network. The design aims to reduce EV charging costs while considering the planned routes, traffic conditions, user preferences, and additional incentives from the energy provider. Similarly, the authors in [Huang 2018] deployed a charging station scheduling algorithm based on the Bitcoin lightning network to minimize various user-incurred costs, waiting time, and distance traveled to the charging point. The work in [Knirsch 2018] resolved the optimal charging station selection problem using the cost-distance trade-off. Although the solution was not implemented practically, both Bitcoin and Ethereum blockchain networks were discussed as a reasonable choice for a small number of EVs. A real-time cryptocurrency-based incentive approach was proposed in [Zhang 2018] to maximize the usage of renewables in the energy system consisting of EV, charging station, PV generation, and battery ESS. Researchers in [Pajic 2018] suggested to couple the Ethereum blockchain with an EV valley-filling charging strategy to reduce overhead on the electric grid. Another example of Ethereum usage was demonstrated in [Liu 2018], where the energy flows of a microgrid consisting of EVs with their charging stations, residential homes, and renewable generation were managed to minimize the impact of injecting or consuming an excessive amount of power into the grid. The authors in [Baza 2019] developed an Ethereum-based EV charging coordination scheme by substituting the PoW consensus mechanism with the PoA.

The second group of works deploys consortium blockchain. Researchers in [Kang 2017] applied an iterative double-auction mechanism to resolve the optimal energy allocation problem. Their system included plug-in hybrid EVs, charging stations, local energy aggregators, and smart meters. The solution aimed to maximize social welfare using consortium blockchain with the PoW consensus mechanism. The authors in [Su 2018] focused on satisfying the needs of EVs while maximizing the operator's utility in the smart community with PV generation and ESS. The formulated energy blockchain system deployed a reputation-based delegated byzantine fault tolerance algorithm to reach consensus among the participating nodes. Yet another consensus scheme, proof-of-reputation, was implemented in [Wang 2019b] to maintain regional energy balance and maximize EV user's utility and satisfaction whilst deploying wireless EV charging technology. The capabilities of EV charging were extended toward V2G and vehicle-to-vehicle concepts in the work of [Huang 2019], where the optimal charging scheduling was considered to maximize EV satisfaction while minimizing the charging costs.

Researchers in [Li 2019] gathered EV manufacturers, EV service providers, EVs, charging station operators, and software providers in a consortium blockchain enabled through Hyperledger Fabric with improved PBFT consensus mechanism. Their work aimed to minimize overall load variance under the distribution network by shifting peak loads while respecting power flow constraints and EV charging demands. The authors in [Fu 2020] included the government as a participant of the consortium blockchain to achieve fair and reasonable EV allocation to charging stations. The proposed EV management scheme addressed the profit allocation problem among various energy companies that provide EV charging services. An incentive economic-based mechanism enabled through EV-coins was suggested in [Chen 2020b] with a particular focus on PV generation. The authors developed a prioritization ranking algorithm to guide the EV charging patterns toward maximizing renewables' utilization. Researchers in [Zhou 2019] considered the internet of EVs and local energy aggregators for the sake of facilitating demand-response measures through consortium blockchain with standard PoW consensus mechanism.

The third group consists of few studies that implemented private blockchain for energy exchange purposes. The authors in [Lasla 2020] deployed a private Ethereum blockchain with PoA consensus mechanism among the network of four entities: EV, energy provider, smart meter, and utility company. The blockchain in their system manages the market auction mechanism and the billing procedure to maximize social welfare. Researchers in [Umoren 2020] deployed a PBFT consensus mechanism in a private blockchain to minimize operation, transportation, and transaction costs in a peer-to-peer trading network of loads and EVs designated as prosumers.

10.3 Methodology

Motivated by the developments mentioned in the literature review, this thesis explores the real-world application of blockchain technology for EV charging management. The research is conducted within the smart community framework introduced in Chapter 7, thus it focuses on the interaction between the smart building, charging station, and EV.

10.3.1 Choice of blockchain

Despite being a powerful and flexible technology, blockchain is not a uniform solution for any problems. Therefore, one has to choose the most suitable blockchain for the problem at hand and customize it if necessary. According to [Di Silvestre 2018b], the following blockchain features have to be considered when choosing a reasonable blockchain implementation: consensus mechanism, speed, permission model, and resilience. The latter signifies the capability of a blockchain-based system to resist to attacks and malicious behaviors. In the current thesis, the aim is to develop a blockchain-enabled energy exchanges framework applicable to real-world scenarios. Therefore, the following criteria are added to the choice of appropriate blockchain implementation:

- **Deployment** follows the blockchain life-cycle from the development, including listing the requirements and actual programming, to the final release of the system into production.
- **Maintainability** refers to the degree of difficulty and effectiveness with which the intended maintainers can modify the blockchain system through updates. The modifications can contain corrections and error handling, system improvements and adaptation.
- **Scalability** signifies the capability of blockchain to handle the growing amount of participants and transactions. Particularly, it is expressed in how fast the blockchain can reach the consensus among nodes and add a new transaction into a block, and how many transactions per second it can process.

This work investigates and compares four popular development frameworks for DApps that support blockchain implementation of smart contracts: AragonOS [Dunkan 2019], Hyperledger Fabric [Androulaki 2018], Energy Web Chain (EWC) [Energy Web Foundation 2019], and Ethereum basic smart contract [Buterin 2014]. The pre-selection of blockchain frameworks and their technical testing were done by HES-SO Valais, a partner in a Digitalization project described in Chapter 7. The in-depth analysis of blockchains' underlying mechanisms and respective advantages and disadvantages was completed in the current thesis. Table 10.1 summarizes the main features used for qualitative comparison among considered frameworks, where the speed is measured in TPS. One has to note that despite planning to switch to Aragon Chain with PoS consensus mechanism in the upcoming future [Izquierdo 2019], the AragonOS currently deploys DApps on Ethereum mainnet. Therefore, the permissioned Hyperledger Fabric is directly compared with permissionless Ethereum and EWC.

Table 10.1 – Qualitative comparison of various blockchain frameworks

Blockchain	Consensus	Speed [TPS]	Permission
Hyperledger Fabric	voting-based	3000	permissioned
Ethereum	lottery-based	15	permissionless
EWC	reputation-based	76	permissionless

As can be seen in Table 10.1, Hyperledger Fabric and Ethereum use different types of consensus algorithms, which are Apache Kafka and PoW, respectively. The lottery-based Ethereum consensus mechanism scales well to a large number of nodes but results in a longer time to finality than Hyperledger Fabric. If two winning miners simultaneously propose a new valid block, the blockchain will experience a temporary fork, and the acceptance of the blockchain state by all nodes will be delayed. The voting-based algorithm of Hyperledger Fabric, instead, provides low-latency finality but scales less efficiently as the time to reach consensus increases with the expansion of the network. Moreover, the algorithm's crash fault-tolerant nature prevents the blockchain from achieving agreement in the presence of faulty

or malicious nodes [Hyperledger 2021]. The PoA reputation-based consensus mechanism of EWC is a hybrid approach that resides in-between the lottery-based and voting-based consensus mechanisms. The PoA relies on a set of trusted miners, called authorities, to take on the leader's responsibility for a new block creation in a rotation manner. Despite being faster and more energy-efficient, the PoA consensus mechanism does not claim a full decentralization while questioning immutability. As authorities' identities are visible to everyone in the network, such things as censorship, blacklisting, and third-party manipulation can be potentially achieved, thus threatening to compromise the safety of the blockchain. In energy-specific Ethereum-based EWC, the largest global energy companies host the validation nodes, thus executing the power to approve the new blocks in a highly regulated energy market. The speed of considered blockchain frameworks varies significantly, as Ethereum is currently processing only 15 TPS, while Hyperledger 3000 TPS. However, researchers in [Gorenflo 2019] have demonstrated recently an upscale of Hyperledger Fabric to 20K TPS, while the Ethereum 2.0 is expected to yield 100K TPS in the upcoming future with the switch to PoS consensus mechanism. The EWC is currently capable of processing around 76 TPS [Energy Web 2020]. However, the Energy Web Foundation organization that launched the EWC repeatedly claims that TPS metric is not suitable to assess the scalability of the blockchain correctly.

To discuss the resilience of considered blockchain frameworks, one has to identify potential cyber threats. The attacks on the blockchain can be grouped into five main categories: blockchain network attacks, user wallet attacks, smart contract attacks, transaction verification mechanism attacks, and mining pool attacks [Katrenko 2020]. The permissioned nature of Hyperledger Fabric adds a layer of security by authorizing access only to a predefined pool of participants. Moreover, the business purpose design of Hyperledger Fabric requires the system to recover from attacks fast without compromising sensitive client data. Researchers in [Chen 2020a] have concluded that blockchains using different programming languages and architectures have different vulnerabilities and are thus susceptible to different types of attacks. The smart contracts of Ethereum written in Solidity language are considered to be more vulnerable than the chaincodes of Hyperledger Fabric programmed in Go. The EWC enterprise-grade blockchain exhibits good resilience, specifically due to the known list of validators that contribute to the overall integrity and security of the network. However, the underlying PoA consensus mechanism is widely criticized for being susceptible to distributed denial-of-service attacks in case of insufficiently large pool of validators.

The energy consumption per transaction was not explicitly included in the decision-making criteria due to difficulties in finding related data for blockchain frameworks beyond Bitcoin and Ethereum. However, a rough estimation of orders of magnitude reported in [Sedlmeir 2020] suggests that EWC's and Hyperledger Fabric's energy consumption reaches 10^3 J and 1 J, respectively. Instead, Ethereum's energy usage is in the order of magnitude 10^8 J [Digiconomist 2021]. Despite Ethereum's current energy intensity, which can be explained by the underlying PoW consensus mechanism, further plans to make it more energy-efficient exist. The transaction cost of various blockchain frameworks was likewise not considered among selection criteria due to being highly volatile and dependent on the network's size and load. Specifically,

Ethereum's cost per transaction has fluctuated in the range 4\$-40\$ since the beginning of 2021 [YCharts 2021], while EWC achieved transaction costs as low as 0.1\$ [EWF 2021]. Although EWC's higher transaction throughput and the absence of PoW bring the costs down, one has to keep in mind that the higher transaction cost of Ethereum is partly justified by the fact that ETH, the network's cryptocurrency, is more valuable. Moreover, growing interest in Ethereum transactions increases the demand for miners, thus driving up the gas prices. Notably, the transaction cost concept does not apply to permissioned blockchains. Therefore, transactions on Hyperledger Fabric tend to be free. Further discussion of blockchain's energy-intensity and transaction costs continues in Section 10.6. Table 10.2 summarizes the assessment of blockchain systems across chosen technical criteria on the scale from 1 to 5, ranging from difficulty to easiness. Notably, the weights of considered criteria differ depending on their importance for the real-world implementation.

Table 10.2 – Assessment comparison of various blockchain frameworks

Blockchain	Deployment	Maintainability	Scalability	
Weights	0.5	0.3	0.2	Total
AragonOS	2	5	3	3.1
Hyperledger Fabric	3	3	4	3.2
Ethereum	5	2	2	3.5
EWC	4	3	3	3.5

The deployment of basic smart contract in Ethereum is considered to be more straightforward than the deployment in Hyperledger Fabric and AragonOS, despite the latter two providing smart contract templates for cloning. Indeed, it consists of two steps successive to writing the contract code itself: compiling the smart contract into bytecode and deploying it by sending an Ethereum transaction without specifying any recipients. All these actions can be performed through Remix IDE. In Hyperledger Fabric, the multi-layered access control framework complicates the process as one has first to install the smart contract on peers and then instantiate it on the channel. However, a single chaincode can define several smart contracts at once, which simplifies the development. The AragonOS has a multi-step deployment procedure with the installation of Node.js runtime environment, Web3 provider to interact with Ethereum, MetaMask to sign transactions, and Aragon command line interface to install the new app. The Ethereum-based EWC offers a comprehensive toolkit with open-source templates of energy-specific digital applications to speed up the development of customized DApps. In addition, the general procedure for deploying a smart contract is similar to the one of Ethereum, where installing the mandatory packages and related development environments is required. However, a preliminary step of setting up the Energy Web Decentralized Operating System is necessary. Moreover, a recent analysis of the EWC network has shown that over 80% of the smart contracts were deployed from only three participating entities [Zwanzger 2020], thus indicating either the lack of interest from general public, the novelty of EWC, or the difficulty to deploy the smart contracts in the energy field.

The AragonOS shows the highest maintainability thanks to easily updating the smart contract to a newer version. Such a feature is available in AragonOS due to the specific design solution: the smart contract's logic is decoupled from its location using proxies. Therefore, developers can fix bugs and push enhancements without changing the address of the smart contract. Similarly, in Hyperledger Fabric, the contracts can be upgraded. However, one has to install the contract with the same name and a different version on all peers before upgrading the smart contract. If the order is reversed, certain peers will lose their ability to participate in the network through endorsing transactions. The smart contracts on Ethereum have the lowest maintainability as they are immutable by default. However, certain approaches to enable upgradeability exist. To release the smart contract update, one can deploy a proxy contract that delegates the execution of methods and functions to the implementation smart contract. Such methodology allows switching the logic contract easily, as users interact only with a proxy contract. However, one has to think of such a maintenance option while the smart contract is still in the design stage and is not yet deployed on the blockchain. The life-cycle of smart contracts on EWC network is supposedly guided by the OpenZeppelin secure smart contract library that gives a possibility to securely destroy and pause the smart contract. However, no comprehensive description of its functionality on the EWC blockchain was found.

The scalability criteria was previously discussed with the qualitative comparison of blockchain frameworks. The Hyperledger Fabric currently scales better than AragonOS, EWC, and Ethereum due to its permissioned nature and absence of the PoW consensus mechanism. However, upcoming releases of Ethereum 2.0 that will potentially include sharding to increase the amount of TPS, and Aragon Chain aim to resolve the scalability issue. For EWC blockchain, the great scalability promises are held due to its PoA consensus mechanism. On a side note, the scalability of AragonOS is seen less of a challenge due to focusing mainly on governance and not on transactions. Particularly, large decentralized organizations can define a quorum to simplify the consensus process, thus eliminating the need to achieve 51% majority of the whole network.

To summarize, the four considered blockchain frameworks for smart contract development vary in their purpose, features, and challenges they are facing. As EWC and Ethereum both acquired the same amount of points in the proposed subjective evaluation scheme, a choice has to be made between the two. Although the prior is specifically suitable for energy-related applications and exhibits lower energy usage and transaction costs, it is not considered a fully decentralized blockchain. Moreover, the upcoming change of Ethereum's consensus mechanism to PoS is expected to significantly reduce Ethereum's energy consumption and drive down transaction costs due to increased transaction throughput. Thus, with the focus on fast deployment and high accessibility of the blockchain network for testing purposes in real-world scenarios, the choice of the basic Ethereum smart contract is concluded as the most appropriate.

10.3.2 Smart contract

The Ethereum smart contract's requirements and code are developed in collaboration with researchers from HES-SO Valais. Written in Solidity programming language, the smart contract is used to settle energy exchanges between charging stations and EVs in the form of blockchain transactions. The following entities with their particular attributes are considered participants of the smart contract:

- **Smart buildings** are the central parties in the smart contract defined by their Ethereum addresses. Each smart building is assumed to possess at least one charging station and eventually one EV, uniquely differentiated by their `idChargingStation` and `idVehicle` identifiers. The charging stations and EVs do not have their own Ethereum addresses as in the charging process they act on behalf of the smart building they belong to. Therefore, a charging transaction is conducted between two smart buildings, where one behaves as the energy provider and the other - as the energy consumer. If the EV is charged at the smart building it belongs to, this smart building takes on a double role resulting in a transaction with itself. Each smart building stores the history of its charging transactions in the `chargingTransactions` list. The smart buildings' respective energy balances are kept in the `buildingBalances` list, which can be queried by the smart building's address. To participate in the charging process, the smart building's registration is required and is verified using the `onlyRegisteredBuildings` modifier.
- The **owner** of the smart contract is the sender of the transaction that deploys this smart contract on the blockchain. The owner is characterized by an Ethereum address and is endowed with two exclusive capabilities enabled by the `onlyowner` modifier: registering new smart buildings and resetting their energy balances. The prior allows the owner to add a new building to the `registeredBuildings` list, thus enabling its participation in energy exchanges between EVs and charging stations. The latter gives the contract owner a possibility to reset the energy balance of a specific smart building to zero if needed. Moreover, the owner has the right to disable a smart building by modifying its status in the `registeredBuildings` list.

The `ChargingTransaction` object of the smart contract is defined as a custom struct containing the state variables presented in Table 10.3:

Table 10.3 – The charging transaction object

Field name	Data type
<code>addressBuildingSupplier</code>	address
<code>addressBuildingConsumer</code>	address
<code>startDate</code>	uint
<code>endDate</code>	uint
<code>energyConsumption</code>	int
<code>idVehicle</code>	uint
<code>idChargingStation</code>	uint

At the beginning of the smart contract, the following public and private mapping functions are defined, where the prior type enables an automated getter-creation in the Solidity language:

- `mapping (address => bool) public registeredBuildings`
allows for a quick verification of the smart building's registration.
- `mapping (address => int) public buildingBalances`
gives a possibility to check the balance of the smart building of interest.
- `mapping (address => ChargingTransaction[]) private chargingTransactions`
associates the list of charging transactions with the smart building's address and allows viewing the charging history from within the smart contract.

The `addChargingTransaction` function is the core of the smart contract. The function uses the `onlyRegisteredBuildings` modifier, thus allowing only the charging stations from registered smart buildings to initiate charging transactions. The input to the function contains the following variables inherited from the `ChargingTransaction` object:

- `address _consumerAddress`
- `uint _startDate`
- `uint _endDate`
- `int _minusEnergyConsumption`
- `int _energyConsumption`
- `uint _idVehicle`
- `uint _idChargingStation`

Notably, the supplier's smart building address is absent from the function input as it is automatically defined by the `msg.sender` property of the charging station issuing the transaction. The `_minusEnergyConsumption` variable differs from the `_energyConsumption` only by its negative sign. As it was explained previously in Section 10.1.1, each operation on Ethereum blockchain is associated with the specific amount of gas that defines its computational complexity. Thus, the introduction of a negative `_minusEnergyConsumption` variable helps to avoid the additional cost related to the subtraction. The `addChargingTransaction` function's body contains the following expressions:

- `chargingTransactions[msg.sender].push()` is used to record the charging transaction on the side of the smart building owning the charging station. The push function appends the `ChargingTransaction` with the amount `_minusEnergyConsumption` to the end of the transaction list.

- `chargingTransactions[_consumerAddress].push()` stores the transaction on the side of the smart building owning the EV. The respective energy amount defined by the `_energyConsumption` variable is positive.
- `buildingBalances[msg.sender] += _minusEnergyConsumption` adjusts the energy balance on the supplier's side.
- `buildingBalances[_consumerAddress] += _energyConsumption` adjusts the energy balance on the consumer's side.

10.4 Implementation

The real-world case study used to validate the proposed blockchain-supported EV charging management approach is conducted within the smart community framework introduced in Chapter 7. The infrastructure of the demonstration site located in Val d'Hérens alpine region was developed and installed by HES-SO Valais in the course of the Digitalization project, previously described in Chapter 7. Therefore, in the testing phase of the proposed framework, the hotels fulfill the role of smart buildings.

10.4.1 Process flow

Figure 10.3 describes the steps of the EV charging process flow. The main entities implicated in the process are the hotel owners and managers, EV charging stations, and hotel guests. The prior participate in the process using the web interface detailed in Section 10.4.2, while the latter issue charging-related commands through the mobile application described in Section 10.4.3. Although both web interface and mobile application were developed by HES-SO Valais, they are presented in the current thesis to provide a complete overview of the functionality and viability of the proposed smart community framework.

To enable the EV's participation in the blockchain-supported charging process, the hotel owner must register the EV in the web interface. Once the EV can be found in the hotel's digital vehicle inventory, this step becomes unnecessary. The registration procedure generates a unique vehicle identifier `idVehicle` together with the username-password pair. The latter is shared with hotel guests to enable access to the mobile application used for managing the EV charging process. Besides the standard way to login, the EV user can utilize the generated QR code, thus avoiding memorizing complex login details. It is expected that the password and QR code will be automatically regenerated.

Once the guest has rented the EV and has connected to the mobile app, the charging process unfolds as follows. First, the EV driver chooses an unoccupied charging station on the map. Second, the user optionally shares the EV's State of Charge (SOC) and the available time to spend at the charging station. If transmitted by the EV driver, these details allow the charging station to optimize the charging process according to the approach suggested in Chapter

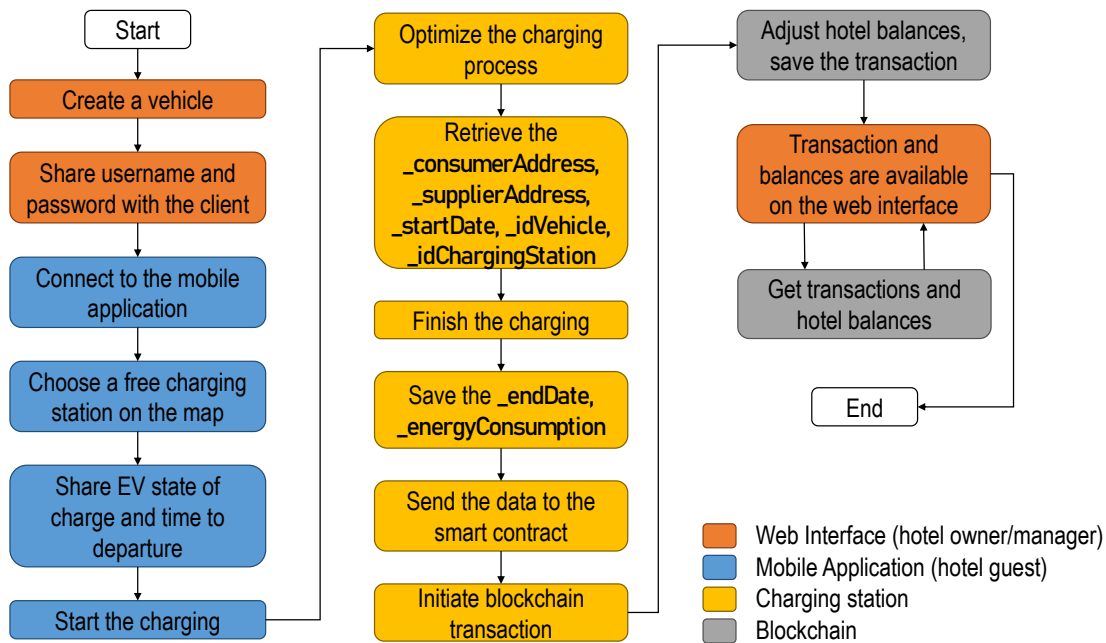


Figure 10.3 – EV charging process flow

9. Third, the EV user starts the charging process by clicking the dedicated button on the mobile application. Otherwise, the charging process can be initiated by the hotel's manager from the web interface. From that moment further, the charging station manages the EV charging process flow. During the EV charging, the charging station collects the data required to populate the `ChargingTransaction` object and execute the `addChargingTransaction` function. The charging station retrieves such data from the web interface using the Application Programming Interface (API). The EV charging ends with either the user terminating the process through the mobile application or the charging station when the maximum SOC is reached. Once the charging is complete, the charging station initiates the blockchain transaction by sending the data to the smart contract. The hotel owner can access the history of charging transactions and their details by interacting with the blockchain.

10.4.2 Web interface

The web interface designed by HES-SO Valais eases the hotel owners' interaction with the blockchain system architecture and simplifies EV management. Once the hotel is connected to the green mobility program and the necessary hardware, such as UniPi, is installed, the hotel owner can create the hotel's account on the web platform. The main dashboard contains various instruments for managing the hotel's participation in the program. Specifically, the three major tabs, namely energy flow, charging station management, and clients, are designed to give the hotel owner an interactive and visual experience of EV management.

Figure 10.4 depicts the energy flow tab, where the hotel owner can monitor the status of its major energy-producing and consuming assets in real-time. Particularly, the tab shows the quantities of PV and grid power consumed by the hotel, the amount of PV fed back to the grid, and the power consumed by the charging station to refill the EV battery.

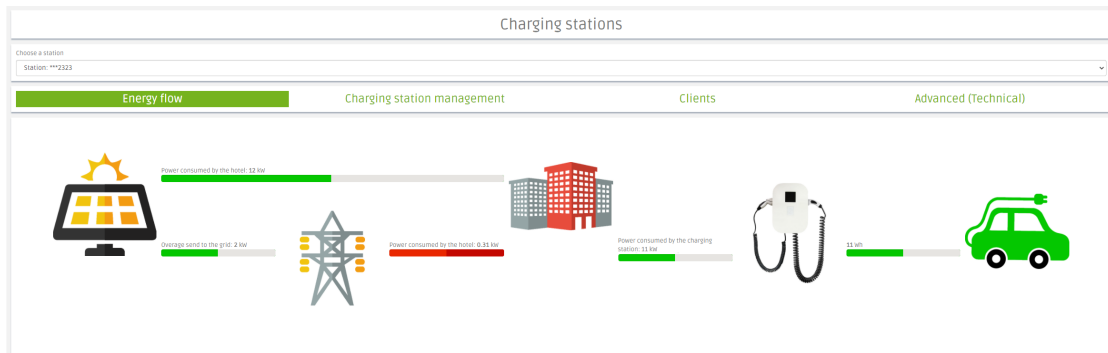


Figure 10.4 – Energy flow tab on the web interface

Figure 10.5 demonstrates the charging station's control panel, where the hotel owner can see the charging station's status and manipulate the charging process. The owner can particularly influence the charging station's availability and switch to the charging process's manual control mode. The latter gives the hotel owner the possibility to set the charging station's desired current, which translates directly to the power output.

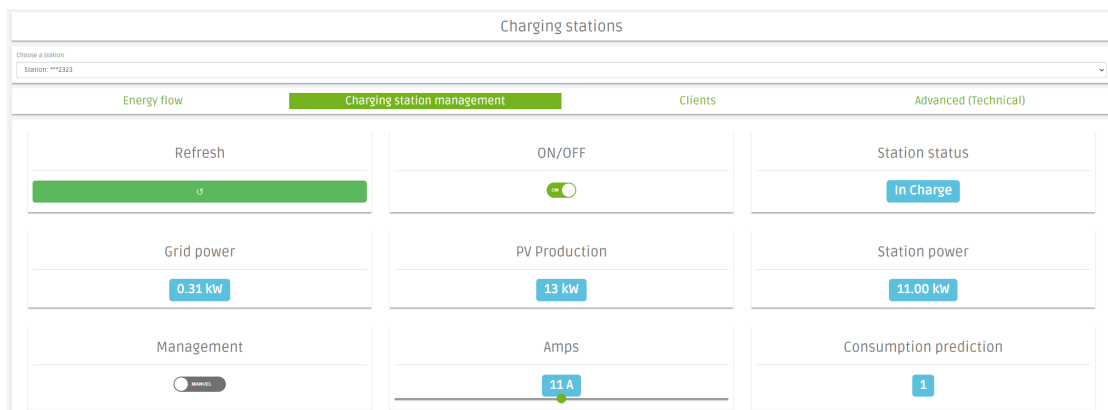


Figure 10.5 – Charging station management tab on the web interface

The clients' tab depicted in Figure 10.6 is the hotel's point of interaction with the vehicle inventory and the blockchain system. The history of charging transactions is displayed on the tab's left-hand side, where the date, identifiers of the EV and the charging station, energy, and charging duration can be seen. The hotel's energy balance is placed underneath the charging transactions and can be used for the hotel's reporting at desired time periods. The tab's right-hand side provides the hotel owner an overview of the EV inventory with the possibility to register additional vehicles.

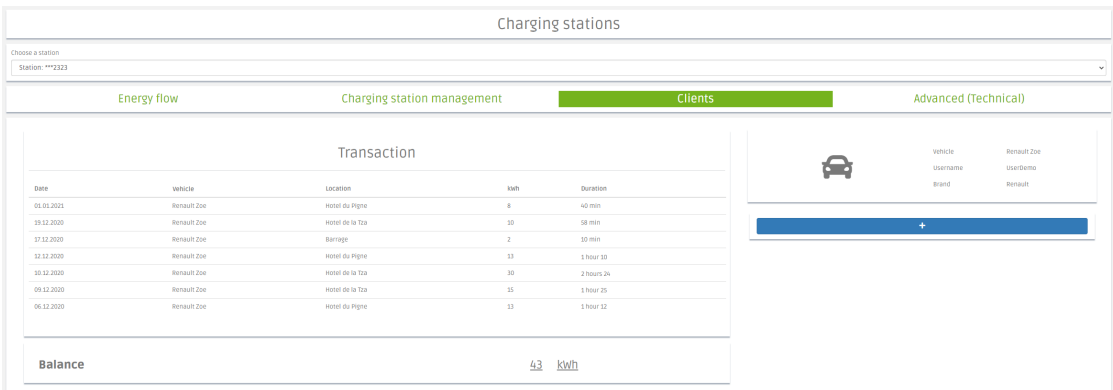


Figure 10.6 – Clients tab on the web interface

10.4.3 Mobile application

The mobile application, equally developed by HES-SO Valais, is the means for the EV user to interact with charging stations participating in the green mobility program. The application developed in NativeScript is available for both mobile operating systems Android and iOS. The following figures demonstrate the mobile application's design and functionalities. Figure 10.7 depicts the login page, where the EV user is asked to either input the username-password pair or scan the QR code. Both are transmitted to the user by the hotel owner or manager at the beginning of EV rental. Figure 10.8 shows the EV user's profile, where the username is displayed. In the future releases of the application, other information about EV, such as model, battery capacity, charger type, etc. will appear on the account screen.

Figure 10.9 presents the map screen that depicts the charging stations in the area and their status. If the charging station is not occupied, the icon's color on the map appears green and red otherwise. Once the EV driver chooses where to charge, the energy feasible route to this charging station can be constructed and displayed by the digital routing energy service proposed in Chapter 8. Clicking on the charging station icon opens additional details, as shown in Figure 10.10. In particular, the user can see the name of the charging station, its address, the hotel it belongs to, the types of charging plugs at disposal, and the charging station's availability. Optionally, to optimize the charging process concerning the hotel's demand, PV production, and EV driver's satisfaction the EV charging digital energy service presented in Chapter 9 can be utilized. Particularly, the EV driver can input the EV's SOC at arrival and the time planned to spend at the charging station. Otherwise, the required charging time is calculated based on the maximum charging power of EV. The charging process is activated using the green button depicted in Figure 10.10.




Figure 10.7 – Login page



Figure 10.8 – Account page

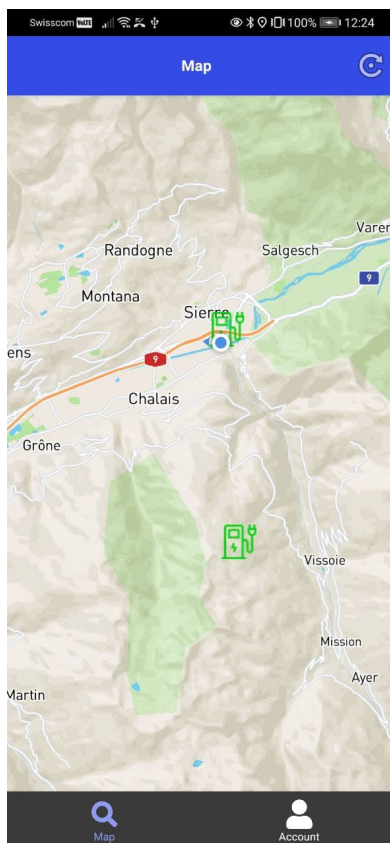


Figure 10.9 – Map of charging stations

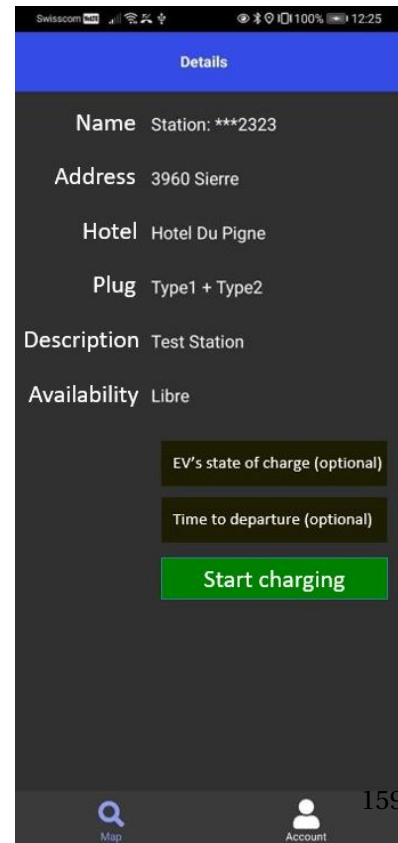


Figure 10.10 – Charging station details

10.5 Case study

To validate the proposed blockchain system architecture and demonstrate its application to the use case of EV charging, the process flow depicted in Figure 10.3 was recreated in a real-world case study by HES-SO Valais. The two hotels, Hotel du Pigne and Hotel Aiguille de la Tza, both located in Val d'Hérens region in Switzerland, act as demonstration participants. The prior hotel owns the EV, while the latter provides the charging station. The experiment is conducted using the UniPi Axon S105 programmable logic computer installed inside the charging station, with the following characteristics: 1.2GHz quad-core ARM CPU, 1GB RAM, 8GB eMMC onboard memory, and 1Gbit Ethernet connection. Figure 10.11 depicts the charging infrastructure used for testing.



Figure 10.11 – Real-world test infrastructure

To enable the vehicle's participation in the blockchain-supported charging, the hotel must register the EV in the inventory using the web interface described in Section 10.4.2. Thus, Hotel du Pigne that already owns a Tesla Model X, uses the 'clients tab' in their web interface's account to add a new EV. The respective dialog window shown in Figure 10.12 prompts the hotel manager to input the necessary information for describing the new EV. The hotel manager makes the following entries to add the Hyundai Kona [EVD 2021c] to the vehicles' inventory and confirms his choices with the blue create button:

- **Username:** Kona
- **Password:** 1234
- **Brand:** Hyundai
- **Model:** Kona
- **Battery size [kWh]:** 64
- **AC board charger [kW]:** 11
- **DC board charger [kW]:** 77
- **Cable type:** 2 (2 = Combined Charging System (CCS), type 2)

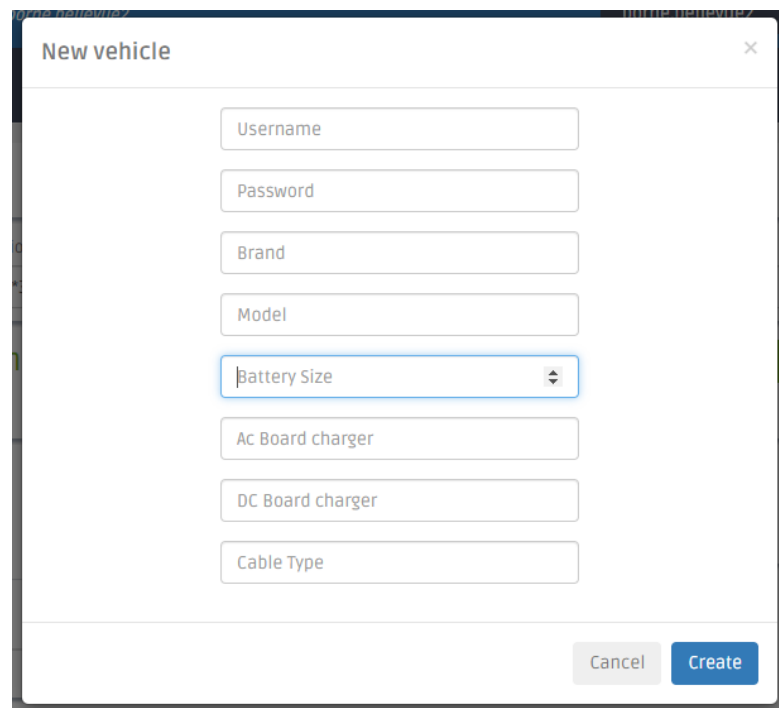


Figure 10.12 – The "Add new vehicle" dialog window

A respective confirmation message depicted in Figure 10.13 pops up on the web interface screen to validate the new addition to the EVs' inventory. The updated list of EVs owned by the Hotel du Pigne is shown in Figure 10.14.

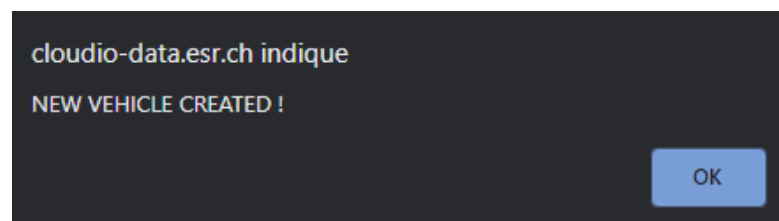


Figure 10.13 – A confirmation message at the registration of a new EV

Once the new EV is registered, a QR-code with its username and password is generated and can be shared with the hotel's guests willing to rent the vehicle. Knowing these details allows the guests to access the mobile application described in Section 10.4.3 and utilize it for charging management. To conduct the charging process, the user logs in to the mobile application using the provided credentials and chooses the charging station of the Hotel Aiguille de la Tza as destination. At the arrival, the EV is charged to 50%, which is optionally reported in the respective field of the screen in Figure 10.10. Omitting the time to departure input, the EV charging process is initiated by clicking on the green 'start charging' button.




	Username	Tesla
	Brand	Tesla
	Model	X
	Battery Size	100
	Username	Kona
	Brand	Hyundai
	Model	Kona
	Battery Size	64
		

Figure 10.14 – The EV inventory of Hotel du Pigne

To form the `ChargingTransaction` object, at the beginning of the charging process, the charging station retrieves the following information:

- `_consumerAddress` 0x150def7979a963fd24ed9b626b612f72343cedab
- `_supplierAddress` 0x3d9c273236233600b98abb4332d12f0a080b3d69
- `_startDate` 1612341928158
- `_idVehicle` 3
- `_idChargingStation` 1

Once the charging is complete, the additional details about the charging process are included:

- `_endDate` 1612342150304
- `_energyConsumption` 0.223 [kWh]

After all the necessary information is collected, the `addChargingTransaction` function implemented in the smart contract is called by the charging station. Once such blockchain transaction is validated, its record can be observed on the Etherscan [Etherscan 2021] as shown in Figure 10.15:

Transactions

Latest 25 from a total of 46 transactions

Txn Hash	Block	Age	From	To	Value	Txn Fee
0xd72fac6c9d7d39e62...	9590517	17 mins ago	0x150def7979a963fd24...	0x69b320f9284183c0e...	0 Ether	0.000686092558

Figure 10.15 – The record of transaction on Etherscan

One has to note that the resulting transaction constitutes the call to the smart contract and is conducted between the hotel possessing the charging station and the smart contract itself. Thus, the hotel owning the EV does not appear as the beneficiary of the transaction due to its public address being included as the input to the `addChargingTransaction` function. Since the transaction is sent automatically by the charging station, the private key of the consumer's Ethereum account is stored locally on the charging station's computer. To ensure the security and protect the system from malicious attacks, the private key can be encrypted to avoid being revealed to unauthorized users. The address of the deployed smart contract is referenced as follows:

- `_contractAddress` 0x69B320F9284183C0E97f21a7956e6D718a62939e

Once the transaction is validated, the charging record and the updated energy balance appear on each hotel's web interface. Therefore, in Figure 10.16 the Hotel Aiguille de la Tza sees the last charging transaction with Hyundai Kona in green, as it was the hotel providing the energy. The same transaction appears in red in Figure 10.17 for Hotel du Pigne, as it was the hotel owning the EV that consumed the energy. Although the sign varies, the total resulting energy balance of 0.15 kWh is the same for both hotels, as they were the only ones included in the testing procedure. Thus, such energy balance signifies the correctness of the accounting and the execution of blockchain transactions.

01/02/2021 14:15:05	00:00:32	Tesla	Hotel Du Pigne	0.016
01/02/2021 14:36:15	00:00:46	Tesla	Hotel Du Pigne	0.023
02/02/2021 08:56:26	00:00:55	Tesla	Hotel Du Pigne	0.027
02/02/2021 09:08:07	00:00:13	Tesla	Hotel Du Pigne	0.007
03/02/2021 09:45:28	00:03:42	Kona	Hotel Du Pigne	0.223

Balance

0.15 kWh

Figure 10.16 – The web interface of the Hotel Aiguille de la Tza

01/02/2021 14:15:05	00:00:32	Tesla	Hotel La Tza	0.016
01/02/2021 14:36:15	00:00:46	Tesla	Hotel La Tza	0.023
02/02/2021 08:56:26	00:00:55	Tesla	Hotel La Tza	0.027
02/02/2021 09:08:07	00:00:13	Tesla	Hotel La Tza	0.007
03/02/2021 09:45:28	00:03:42	Kona	Hotel La Tza	0.223
Balance -0.15 kWh				

Figure 10.17 – The web interface of the Hotel du Pigne

10.6 Results

To assess the performance of the proposed blockchain-supported EV charging framework, the following metrics widely applied in the blockchain literature are discussed:

- **Transaction latency** or average transaction time is the time elapsed between the transaction's generation and its final appearance in the block on the blockchain. Despite being widely used, this metric varies strongly depending on the following parameters: the number of simultaneous transactions on the network, the average gas price of every pending transaction, and the gas price the user is willing to pay. If the network is overloaded, users will have to set a higher gas price for their transaction to be processed and written on the blockchain by miners.
- **Transaction fee** is a cryptocurrency fee collected from users to process the transaction on the blockchain network. The fee differs depending on the complexity of the transaction, the gas price set by the user, and the price of Ethereum at the date of the transaction. This metric is related to the average transaction time since a higher transaction fee results in a shorter transaction time.
- The **contract deployment fee** is the price of the smart contract's initial deployment on the main Ethereum blockchain. This must be done once at every update of the smart contract's code. The fee works the same way as the regular transaction fee, and the price depends on the same parameters.
- The **number of nodes** is a good measure to understand the size of the blockchain network. However, it is more suitable for private blockchains. As the current work was conducted using Ethereum public network, the total number of nodes, which at the moment of writing this thesis equals 12473 on Ethereum, is not a metric of interest [The Ethereum Network & Node Explorer 2021].

- **Transaction throughput** is the number of transactions per second that the network can process. Therefore, it gives a good idea of how scalable the system could be in the future. However, such a metric is not applicable to evaluate the suggested methodology's performance as it utilizes Ethereum public blockchain. Thus, the performance is limited by the mainnet's capabilities without having the means to influence it.

To gather the data for analyzing the aforementioned metrics in the current thesis, manual tests were performed on the blockchain by HES-SO Valais. At the time of this work, the price of a single transaction of the designed smart contract on the mainnet was over 80\$. Therefore, for financial reasons, the tests were only conducted on the Ropsten testnet. First, the appropriate gas price had to be set to indicate the desired speed of mining the transaction. The online service [Ethereum 2021] provides the statistics of recommended gas prices based on supply and demand for the computational power of the network needed to process smart contracts and other transactions, as seen in Figure 10.18. At the time of the experiment, the recommended gas price for a standard transaction was 155 Gwei, equal to 0.000000155 ETH.




Figure 10.18 – The recommended gas price as of February 4th, 2021

Once the gas price has been defined, the gas limit had to be set to ensure that the transaction can be executed. If the computational power needed to execute the transaction exceeds the predefined gas limit, the transaction will be aborted as it runs out of gas. In the current implementation of the smart contract, a re-issue of the aborted transaction with augmented gas limit is not foreseen. Thus, the owner of the charging station will have to manually re-send the transaction to the blockchain. However, as the unused gas is reimbursed by the blockchain network, the gas limits are largely overestimated in practice to ensure the transaction's safe validation. As a result, the events of transaction abortion are rare, which can justify the absence of transaction re-emission functionality in the smart contract. Figure 10.19 refers to the estimation provided by the same online service [Ethereum 2021] that predicts transaction cost and mean confirmation time.

Table 10.4 summarizes the experiment's results, particularly reporting the transaction times. The start and end of the transaction are specified as Unix timestamp. Considering a constant gas price was applied, the difference in transaction times can be explained by the number of simultaneous transactions in the network. Due to each Ethereum block having a maximum size of 1.5 million gas, the number of standard transactions that can fit inside is around 70. Therefore, if the transaction is processed close to the end of the block's formation, the transaction time can be fast. Otherwise, the transaction has to wait for the block's completion to be recorded on the blockchain.

Transaction Inputs

Gas Used*

Gas Price*
☐ Fastest (181 Gwei)
☐ Fast (160 Gwei)
☐ Average (131 Gwei)
☐ Cheap (131 Gwei)
☒ Other

Predictions: Gas Used = 330000; Gas Price = 155 gwei

Outcome	
% of last 200 blocks accepting this gas price	80
Transactions At or Above in Current Txpool	261
Mean Time to Confirm (Blocks)	3.3
Mean Time to Confirm (Seconds)	53
Transaction fee (ETH)	0.05115
Transaction fee (Fiat)	\$84.34635

Figure 10.19 – The gas calculator

Table 10.4 – The data collected during the experiment

Num	Start	End	Time [ms]
1	1612434344629	1612434367000	22371
2	1612434037000	1612434042535	5535
3	1612433871000	1612433878263	7263
4	1612433566000	1612433568287	2287
5	1612433265875	1612433297000	31125

The following metrics are derived from the conducted experiment:

- **Transaction latency:** 13.716 [s]
- **Transaction fee:** 0.00074103 [ETH] \approx 1 [CHF]

However, all transactions on the Ropsten testnet are worthless and do not reflect the price that would have been paid on the main Ethereum network. Nevertheless, the following discussion of the proposed framework's viability considers the calculated transaction fee as the reference.

The typical EV charging business model usually involves three parties who collectively determine the price of EV charging, with their respective roles, responsibilities, and cost structures defined as follows [Verbeek 2020]:

- **Charging station owner**, who owns the charging station and respective location. The type of ownership can be semi-public, public, and private.
- **Charging Station Operator (CSO)**, who is responsible for the management, maintenance, and operation of the charging station.

- **Electric-Mobility Service Provider (EMSP)**, who offers EV charging services to end-customers.

According to [Madina 2016], the cost structures of CSO and EMSP are defined in Equations 10.1 and 10.2, respectively, where C_i is the infrastructure cost related to annual amortization of charging station investment and operation and maintenance, C_e is the electricity bill, C_c is the communication cost, C_{mp} is the access to marketplace cost, C_o is the staff and overhead cost, and C_{cm} is the customer management cost:

$$C_{CSO} = C_i + C_e + C_c + C_{mp} + C_o \quad (10.1)$$

$$C_{EMSP} = C_c + C_{mp} + C_o + C_{cm} \quad (10.2)$$

One has to note that Equation 10.2 reflects only the EMSP's fixed costs, while the variable costs are assumed to be passed through directly to EV users [Madina 2016]. Moreover, the terms in Equations 10.1 and 10.2 are used to indicate the nature of compounded costs without implying any similarity in orders of magnitude. To generate revenues and effectively compensate for their expenditures, the CSO determines the minimum average charging price, which can be both energy-based [CHF/kWh] and time-based [CHF/min], and EMSP defines the minimum value of a subscription fee, which allows the EV driver to access the EMSP's charging network [Madina 2016]. Importantly, if the EV charging process is conducted outside the EMSP's network, the roaming fees must be paid. At the end of the charging event, the EV driver pays to EMSP, who shares its revenue with CSO, who in turn compensates the charging station owner.

As one can notice, utilization of conventional EV charging business model leads to the accumulation of various non-charging related costs, such as C_{mp} , C_o , C_c , C_{cm} , and roaming fees. Implementing the proposed blockchain-supported EV charging management framework can potentially reduce these costs by automating the charging records procedure, impacting the role of EMSPs, and influencing the rules to access the EV charging market. To test this hypothesis, the example of EVPass, one of the largest public EMSPs in the highly fragmented Swiss market [Open Charge Map 2021], is considered. The driving-related statistics required to assess the conventional business model are summarized in Table 10.5 along with technical details of three popular EV models: Citroen C-Zero (CCZ), Hyundai Kona (HK), and Tesla Model X Long Range (TMX). In this study, it is assumed that one charging event performs complete battery charging from 20% to 100%, although in practice, EV drivers prefer to top up their batteries before the minimum SOC is reached. The number of annual recharges N_i and annual energy consumption E_i are calculated according to Equations 10.3 and 10.4, where i is the EV model:

$$N_i = \frac{D_{av} * 365}{(100\% - SOC_{min}) * D_i} \quad (10.3)$$

$$E_i = N_i * (100\% - SOC_{min}) * B_i \quad (10.4)$$

Table 10.5 – Case study data

Variable	Notation	Value	Unit
Minimum EV SOC	SOC_{min}	20	%
HK battery size [EVD 2021c]	B_{HK}	64	kWh
HK real driving range [EVD 2021c]	D_{HK}	395	km
CCZ battery size [EVD 2021b]	B_{CCZ}	16	kWh
CCZ real driving range [EVD 2021b]	D_{CCZ}	85	km
TMX battery size [EVD 2021b]	B_{TMX}	100	kWh
TMX real driving range [EVD 2021b]	D_{TMX}	440	km
Average daily distance driven in Switzerland by car [OFS 2015]	D_{av}	23.8	km
Number of annual recharges of HK	N_{HK}	28	-
Annual energy consumption of HK	E_{HK}	1434	kWh
Number of annual recharges of CCZ	N_{CCZ}	128	-
Annual energy consumption of CCZ	E_{CCZ}	1639	kWh
Number of annual recharges of TMX	N_{TMX}	25	-
Annual energy consumption of TMX	E_{TMX}	2000	kWh

The EVPass EMSP offers three different subscription services along with Pay As You Go (PAYG) option. The Night, Day, and Anytime subscription packages amount to 550 CHF, 990 CHF, and 1320 CHF annually, while the first PAYG option charges a 1.5 CHF flat-rate for each charging event along with 0.5 CHF/kWh, and the second PAYG option charges 59 CHF annually along with 0.45 CHF/KWh [EV Pass 2021]. Single charging event-related cost for EV driver per EV model under different payment options is calculated using the number of annual recharges and annual energy consumption determined in Table 10.5 and is presented in Table 10.6. According to research conducted in [Krug 2021], the division of revenues in the EV charging market can be approximated as follows: charging station owner 17.6%, CSO 21.2%, EMSP 7.3%, and electricity provider 53.9%. Therefore, assuming that the blockchain-supported charging system can potentially impact the role of EMSPs and reduce some of the expenses mentioned earlier, 7.3% of the charging-event costs calculated in Table 10.6 could serve as the reference to set cost-efficiency objectives for blockchain in EV charging.

Table 10.6 – Charging-event related costs for EV driver in [CHF]

EV Model	PAYG 1		PAYG 2		Night		Day		Anytime	
	Total	EMSP	Total	EMSP	Total	EMSP	Total	EMSP	Total	EMSP
HK	27.1	2.0	25.1	1.8	19.6	1.4	35.3	2.6	47.1	3.4
CCZ	7.9	0.6	6.2	0.5	4.3	0.3	7.7	0.6	10.3	0.7
TMX	41.5	3.0	38.4	2.8	22.0	1.6	39.6	2.9	52.8	3.8

As shown in Table 10.6, the potential impact of implementing blockchain in the EV charging process on cost reduction varies depending on the EV model and becomes more pronounced the lower the amount of annual recharges is. Particularly, in the case of CCZ, the blockchain-incurred costs would have to demonstrate higher cost-competitiveness than for HK and TMX to offset the EMSP's expenditures effectively. If the previously calculated transaction fee is taken as the reference value, the EMSP's revenue per charging event of CCZ appears to be inferior to 1 CHF, thus resulting in no cost reductions. However, for TMX and HK, the average amount of blockchain-related savings in a single charging event can be estimated as 1.82 CHF and 1.24 CHF, respectively. Thus, the average between the two is 1.53 CHF. However, the cost advantage of blockchain solution in Night, Day, and Anytime subscription models would disappear if the annual amount of charging events rises to 41, 73, and 97 times, respectively. Therefore, the value of transaction's fee on the blockchain network is an important parameter to be considered. Moreover, to ensure the blockchain's success in EV charging, one can deliberately define the target transaction fee and aim to achieve it by choosing the blockchain type with appropriate characteristics. Considering the future development of electric mobility, the current amount of 67184 EVs on Swiss roads [OFS 2020] is supposed to increase substantially. To estimate the amount of annual country-wide potential blockchain-related savings, it is assumed that Swiss EV stock will maintain its share in the European market, which by 2030 is projected to increase from current 3M EVs [ACEA 2020] to 30M [Abnett 2020], 140M [IEA 2020a], and 245M [IEA 2020a] in conservative, moderate, and optimistic scenarios, respectively. Thus, the estimated amount of EVs on Swiss roads in 2030 varies from a conservative 700K to an optimistic 1.7M. Considering previously estimated blockchain-related average savings of 1.53 CHF per charging event and assuming 30 annual recharges to account for more frequent recharging behavior of EV drivers in practice, the collective financial savings would roughly constitute 32M CHF and 78M CHF in conservative and optimistic scenarios, respectively. Devoting such savings to the installation of new charging stations would strengthen the Swiss charging landscape that currently amounts to 6600 charging points [Statista 2020].

Besides financial and security advantages, implementing the proposed blockchain-supported EV charging framework would result in an additional set of benefits facilitating EVs' adoption. First, simplified access to the market would give private charging station owners the possibility to offer their charging assets to a wider public, thus increasing the size and efficiency of the EV charging network. Additionally, blockchain records' immutability would serve as a guarantee for private owners in a trustless and often insecure environment. Second, blockchain implementation would improve the EV driver's experience by eliminating the need to manage several charging subscriptions, choosing only permitted charging stations, and paying roaming fees when charging outside the chosen network. However, one must keep in mind that using blockchain in EV charging processes is currently subjected to certain obstacles. First, it is not fully compatible with EU's General Data Protection Regulation (GDPR) enforced in 2018. Distributed data storage and processing complicate the attribution of responsibility to a single person as requested by GDPR. Moreover, blockchain's fundamental immutability principle contradicts the GDPR's obligation to make data editable and erasable. Therefore,

despite the belief that data encryption can solve the GDPR compliance problem, there is still a long way to achieve full legal compatibility. Second, the modern blockchain's energy intensity and carbon footprint question the appropriateness of adopting such technology for energy management purposes. Although Ethereum's energy consumption of 75 kWh per transaction is significantly less than 1 MWh value reported for Bitcoin [Statista 2021], it still represents a substantial amount. If Ethereum would be a country, its annual electrical energy demand equal to 29.5 TWh would be comparable to the one of Slovakia [Digiconomist 2021]. However, Ethereum's plan to switch to the PoS consensus mechanism in late 2021 is expected to increase energy efficiency and significantly reduce electrical energy consumption.

10.7 Conclusion

To conclude, this chapter proposed a blockchain-supported EV charging management framework that aids the participating entities to monitor and control the charging process and ensures the correct accounting of the energy flows. Constituting the core of the suggested methodology, the designed EV charging-specific smart contract governs respective energy exchanges while recording the charging information in blockchain transactions. The real-world case study conducted in Switzerland has shown the viability of the proposed solution, while assessing the framework's performance according to common blockchain metrics has confirmed the suitability for real-world application. Moreover, the discussion focused on the current EV charging business model has highlighted potential cost-reduction benefits related to blockchain implementation alongside other likely advantages facilitating the adoption of EVs. Thus, the proposed solution, without any doubt, can be utilized as a standalone digital energy service, particularly taking into account the developed web and mobile interfaces. However, even more valuable, the EV charging blockchain-supported digital service efficiently encompasses the routing and charging control applications introduced in Chapters 8 and 9, respectively. The future work to improve the suggested EV charging management framework and enhance the application of blockchain solutions in the field should be conducted in three main directions:

- To ensure the protected use of the blockchain-supported framework and avoid unexpected future situations that can potentially compromise participating entities' safety, the system's security should be tested by conducting simulated experiences of potential cyber threats. Specifically, one should assess the reliability of storing the consumer's private key on the charging station's computer and test various encryption methods.
- To effectively accommodate the future features of a growing and increasingly complexifying transport system, the framework's scalability limits should be analyzed through a set of extensive experiments involving multiple EVs and charging stations, where the charging processes are handled simultaneously. Thus, the respective blockchain performance-related metrics should be reassessed for several scalability scenarios, whereas future developments of the blockchain technology should be anticipated.

- To unleash the energy blockchain's market potential and provide long-term benefits to participating stakeholders, the blockchain's legal status should be unequivocally determined. Although the cryptocurrency aspect is currently not included in the suggested EV charging framework, one can imagine enabling blockchain-supported payments in the nearest future. Therefore, to effectively decouple from the banking sector and embrace decentralization at its fullest, the market rules governing the new blockchain-enabled energy sector should be defined based on principles of transparency and fairness.

Conclusion

Summary

In challenging times of climate change, the large-scale transformation of the energy industry fostered by the energy transition questions the well-established order of energy systems. Thus, to reinvent themselves and fit seamlessly into the new decarbonized and decentralized energy paradigm, the energy industry's stakeholders turn to the merits of digitalization, strongly supported by the exponential development of its key driving forces. To aid utilities and end-users in finding their place in the new energy world and ensure equal long-term benefits for all, concrete actions that embrace the inevitable change brought by the rapidly developing technological landscape have to be orchestrated. Therefore, this thesis proposes a set of frameworks, tools, and models that can effectively boost energy systems' digitalization while creating the basis for the future full-scale industry-wide digital transformation.

Whilst the overall thesis objective was to identify, develop, implement, and evaluate promising digitalization use cases in the electric power sector, Part I of the current thesis focused on the digital energy services applicable for intelligent energy management within smart buildings. Particularly, Chapter 3 provided a decision-making framework that helps building managers and occupants to understand respective energy consumption patterns, detect best-suited windows of opportunities to implement actions improving the energy efficiency, and streamline energy-related operations to unlock the buildings' vast potential for energy savings.

As such, Chapter 4 proposed a novel low-sampling-rate unsupervised load disaggregation algorithm based on activity-chain reconstruction through the Markov process that allows to efficiently breakdown the whole-house aggregated energy consumption into eight distinct categories of household appliances. The algorithm's benchmarking to state-of-the-art methods in the field proved that it achieves similar disaggregation accuracy despite its unsupervised nature while exhibiting lower computational cost. The algorithm's performance was showcased on real-world datasets from Norway and Germany, highlighting the importance of considering regional characteristics, variations in household equipment, behavioral patterns, and consumer preferences to implement disaggregation practices successfully. Chapter 5 suggested two highly-accurate context-based approaches to forecast binary occupancy using advanced ML techniques. While the supervised method predicts occupancy from electrical energy consumption data and the unsupervised one uses ambient environment measure-

ments, both approaches demonstrated prediction accuracies beyond 97%. The algorithms' performance was evaluated on a real-world dataset collected in Portugal. The importance of carefully considering the needs of the targeted digital energy service to understand which accuracy gains and which computational speeds would make a difference in the success of occupancy-centered energy management strategies was emphasized. Chapter 6 proposed a novel occupancy-centric rule-based automation algorithm for HVAC that generates intelligent day-ahead ON/OFF schedules aiming to deliver energy savings without compromising the occupants' thermal comfort. Evaluating the algorithm's performance on two real-world datasets collected in Portugal highlighted the higher energy-saving potential of cooling as opposed to heating. Furthermore, the significance of achieved potential energy savings and the need to plan in advance for the availability of appropriate automation instruments were pointed out.

Part II of this thesis zoomed out from the smart building framework and investigated the impact of digitalization phenomena on the broader energy ecosystem at the level of the smart community. Particularly, it focused on electric mobility as one of the most promising examples of sector coupling that facilitates the integration of utility-scale renewables and increases the overall energy system's efficiency, flexibility, and reliability. The proposed in Chapter 7 smart community ICT-based framework brings together physical and digital assets, thus creating the environment for exploring various digital energy services.

Particularly, Chapter 8 introduced an EV-specific routing method based on RL that aims to generate energy feasible paths from starting point to destination while considering both recharging possibilities at intermediary charging stations and the ability of EVs to recuperate energy. Formulated in the graph-theoretical context, the suggested approach was validated on a segment of a real-world road network located in Switzerland, thus confirming the solution's capability to either constitute a standalone digital energy service or become an auxiliary instrument in broader EV charging management. Chapter 9 developed an extensive methodology based on Markov decision processes to enable deep RL control of EV charging with the objective to simultaneously maximize PV self-consumption and EV's state of charge at departure. Differing by the type of action space used, the three suggested deep RL algorithms were benchmarked against popular control methods in the field, such as rule-based and model predictive control, proving the RL's superior computational performance and better fitness for future mobility systems. Designing such innovative control methods for EV charging paves the way for unlocking the hidden flexibility potential, extending the scope of energy management objectives, and scaling the energy systems to integrate the growing number of emerging energy actors efficiently. Finally, Chapter 10 described a blockchain-based EV charging digital energy service that enables secure and reliable accounting of energy exchanges within the smart community through a specifically designed smart contract. Tightly interlinked with physical infrastructure and implemented in a real-world demonstration site in Switzerland, the suggested service demonstrated blockchain's potential for reducing costs associated with the charging process, innovating the EV charging business model, and transforming the role of key market players, all while facilitating the large-scale deployment of electric mobility.

Perspectives

Aiming to provide a reliable foundation for further digitalization of the electric power sector, the developed smart building ICT-based platform and encompassing it broader smart community framework exhibit a modular structure, highly efficient in both research and industrial contexts. Particularly, the inherent modularity of the suggested frameworks facilitates the digitalization of energy systems in three ways:

- First, it simplifies the development of new complementary digital energy services that can benefit largely from existing information flows, a structured approach to data collection and storage, and considered metering and sensing devices. Therefore, such synergy is expected to accelerate the prototyping, testing, and implementation of new digital applications, reducing the time to market.
- Second, the possibility to select specific digital energy services that are relevant for certain customer groups enables greater personalization, thus increasing the adoption rate of digitalization in energy systems, improving utilities' relationships with their clients, and helping energy providers to transform themselves into trusted energy advisors to fit the developing landscape of a new energy economy.
- Finally, it gives utilities the upper hand in managing critical situations with instantaneous feedback, faster responsiveness, and better customer care. Furthermore, the inherent distinct division of responsibility areas helps to optimize required maintenance processes while minimizing users' inconvenience.

The suggested digital energy services represent versatile applications that can be modified, augmented, and personalized to account for various desired levels of detail, incorporate a greater number of implicated energy actors, and achieve even higher energy efficiencies thanks to ever-growing amounts of data. Notably, one can adapt the developed digital energy services to higher complexities of considered use cases without significant changes in the methodology. Such adaptability is foreseen in the applications at the design stage, where additional specifications of the examined energy systems can be included using features. Besides, a more holistic approach to data collection, treatment, and management should be adopted, enabling greater generalization of proposed methods and giving interested stakeholders the possibility to deploy digital energy services of choice in a trouble-free manner. While further research is required to internalize a common practice of using data-driven techniques for solving the energy industry's most prominent problems, the underlying methodologies presented in the current thesis merely aim to unveil the capabilities of ICT and facilitate their adoption in the industry. Such purpose is supported by the fact that this thesis demonstrates the application of certain ML algorithms and blockchain technology, which, to the best of our knowledge, were not previously deployed in the considered energy use cases.

Conclusion

Beyond improving technical implementation and scalability of the developed methods, the future work to digitalize the electric power sector should be conducted alongside economic, political, and societal dimensions determined in the research framework presented in Chapter 2. Particularly, one should dedicate the efforts to redesign the energy market space, taking into account novel tariff structures, the growing importance of peer-to-peer transactions, and the energy's transformation from commodity to a service. Besides, a thorough cost-benefit analysis of digitalization should be performed on a case-by-case basis, thus providing stakeholders with decision-making guidelines and offering the baseline for assessing digitalization projects in the future. Moreover, new flexible policy frameworks that facilitate further industry digitalization have to be established. Finally, the influence of energy digitalization on creating new societal norms and behavioral patterns has to be studied to ensure that it contributes to society's development without exacerbating inequality and other pre-existing problems. Although digitalization was proven to significantly increase energy efficiency, recent studies demonstrating a positive correlation between ICT deployment and energy consumption's growth raise concerns on the expediency of ubiquitous digitalization [Lange 2020]. Therefore, one has to focus on investigating the environmental sustainability of ICT, paying special attention to operational energy demands of communication networks, new technologies like blockchain, end-user devices, and data centers.

Recommendations

Future development of energy systems is driven by clear technological trends, such as renewables' growth, expansion of electric mobility, and transition towards smart grids. Thus, utilities find themselves in a rather predictable context, where concrete actions taken today will result in foreseeable, understandable, and measurable consequences in prospect. Although the ICT foundation of digitalization can be already characterized as solid since many key technologies are developing at an accelerated pace, the mindsets of major energy players and policymakers still require a change to accept the inevitability of energy systems' digital transformation and steer it onto an efficient, secure, and sustainable path.

Since utilities can expect approximately 20% efficiency gains¹ from AI in the next 1-5 years [Henzelmann 2018], they should dedicate their present-day efforts to building digital expertise through scouting digital talent, enhancing intra-industrial and cross-industrial relationships to efficiently exchange knowledge and experience, and setting up comprehensive data management systems that follow data throughout its full life cycle. The latter, responsible for the collection, storage, processing, usage, and destruction of data, will become the crucial element determining not only the competitive advantage of each individual utility company but also the success of digitalization in general. Furthermore, well-defined data management strategies can improve utilities' relationships with their customers by placing privacy and security at the heart of operations and reduce the risk of potentially dangerous cyber threats through a clear understanding of existing vulnerabilities.

¹Evaluated based on a survey about AI pilots in trading, operations, maintenance, and support functions

On the policy side, one should focus the attention on ensuring appropriate access to timely, robust, and verifiable data [IEA 2017] and revisit existing legislation to incorporate more flexibility for accommodating actual and emerging technologies. Although General Data Protection Regulation (GDPR) can be seen as a great achievement setting the foundation of data governance and protection in the EU, the reality shows that it rather hampers digital innovations than supports them [Voss 2021]. Particularly, certain provisions on data ownership hinder the development and implementation of AI and blockchain, despite these technologies being around long before enforcing the GDPR. Thus, policymakers should find a healthy balance between regulating the field and giving the necessary freedom to advance digitalization even further. At the same time, the technology providers should dedicate efforts to make their developments GDPR-compliant, especially if they want to achieve high acceptance rates among the general public. In addition to issues related to data governance and usage, one should find the answers to very delicate ethics questions. Since we are rapidly progressing towards the moment when sophisticated ML algorithms will make major operational decisions, the fundamental principles of responsibility should be redefined. Particularly, it is relevant for the energy industry, as power outages and blackouts occurring due to incorrect decisions taken by AI can have very tangible, costly, and even disastrous consequences. Thus, specific guidelines need to be established to ascribe responsibility for the wrongdoing of AI.

Concluding remarks

To conclude, if the energy industry's stakeholders will embrace digitalization and commit their efforts towards changing their mindset, strategy, and operations, we might find ourselves in a dramatically different energy world. In a world where power generation, transmission, and distribution are automatically optimized and supply and demand are balanced by AI without any human intervention. In a world where we will never have to see an electricity bill again since all payments will be securely and reliably settled throughout blockchain at a speed that only machines can tackle. In a world where energy management systems in our smart homes will constantly strive for energy efficiency while continuously learning our preferences so that thermostats are always kept at comfortable temperatures and EVs are charged whenever we need them. In a world where digitally enhanced energy systems will be the foundation of our planet's bright and sustainable future for many generations ahead.

A Appendix

A.1 Characteristics of households

Table A.1 – Characteristics and appliances' inventory of Norwegian demonstration households

	House	Apartment
Number of people	4	2
Number of kids	2	0
Person age	42	33
Employment adults	full-time	full-time
Employment kids	student	n/a
TV	2	1
PC & laptop	6	4
HiFi	2	1
Gaming console	2	0
Washing machine	1	1
Dishwasher	1	1
Tumble dryer	1	0
Fridge with freezer	0	1
Fridge	1	0
Freezer	1	0
Kettle	1	1
Coffee machine	1	0
Printer	1	0
Modem internet	1	1
Microwave	1	0
Mobile phone	4	2
Tablet	2	2
Sauna	1	0
Hairdryer	1	1
TV box	1	1
Oven	2	1
Vacuum cleaner	1	1
Fridge room	1	0

A.1. Characteristics of households

Table A.2 – Characteristics and appliances' inventory of German demonstration households

	ID1002	ID1004	ID1007	ID1009	ID1011	ID1012	ID1013	ID1014	ID1017	ID1022
Number of people	2	3	5	2	2	3	2	5	4	2
Number of kids	0	0	2	0	0	2	0	1	1	0
Employment adults	FT	FT	FT	FT	FT	PT	FT	FT	FT	FT
Employment kids	n/a	n/a	ST	n/a	n/a	ST,CH	n/a	ST	ST	n/a
Dishwasher	1	1	1	0	1	1	1	1	1	1
Tumble dryer	1	1	0	1	1	1	0	1	1	1
Fridge with freezer	0	1	1	1	1	1	1	1	1	0
Freezer	1	1	0	1	1	1	0	1	0	1
Coffee machine	1	1	0	0	1	1	1	1	1	1
Printer	1	1	1	1	1	1	1	1	1	0
Microwave	1	1	1	1	1	1	1	1	1	1
Kettle	1	1	1	1	1	1	1	1	1	1
Sauna	1	0	0	1	1	0	0	0	0	0
DVD	0	0	0	0	0	0	0	0	0	0
Aquarium	0	0	0	0	1	0	1	0	0	0
Laptop	1	0	0	1	1	1	1	1	1	0
Fridge	1	1	0	0	0	0	1	0	0	1
TV	1	3	1	3	1	3	1	5	3	2
PC	0	1	1	4	1	1	0	1	2	0
HiFi	1	2	1	1	1	1	1	2	2	2
Gaming console	0	1	0	1	1	0	0	2	1	1

FT = Full-time employment

PT = Part-time employment

ST = Student

CH = Child (unemployed)

Common devices such as washing machine, modem internet, mobile phone, hairdryer, oven, and vacuum cleaner are present in all demonstration households.

The age of each household's principal person is in the range of 30-64 years.

A.2 Characteristics of occupancy detection sensors

Table A.3 – Characterization of sensor technologies used for occupancy detection

Technology	Working principle	Functionality			Cost
		Detect	Count	Track	
Passive Infrared Sensor (PIR) ^{1,2,9}	Detects infrared radiation changes caused by the movement of subjects.	✓	✓	x	Low
Smart meter ^{1,2}	Occupancy can be inferred analyzing the changes in magnitude of electricity consumption.	✓	x	x	n/a
Bluetooth Low Power iBeacon (BLE) ^{1,2}	Estimates the number of smart-phones connected and classifies occupants' location based on strength of signal received.	✓	✓	✓	Medium
Chair sensor ^{1,2}	Thresholding method to detect the occupied states of a chair.	✓	✓	✓	High
Vibration sensor ^{1,2,9}	Detects footsteps	✓	✓	✓	High
Ultrasonic sensor ^{1,2,9}	Emits ultrasound pulses to detect occupants' movement. Any interruption in the transmitted pulse corresponds to detected occupancy.	✓	x	x	High
Radio-Frequency Identification Tag (RFID) ^{1,2,9}	Uses radio waves transmission to localize occupants.	✓	✓	✓	Low
CO ₂ sensor ^{1,2,3,9}	Changes of indoor CO ₂ concentration are used for dynamic detection of occupancy.	✓	✓	x	Medium
Video camera ^{1,2,4,9}	Dynamic visual information is used to analyze occupancy.	✓	✓	✓	Low
WiFi ^{1,2}	Detects the number of devices connected to the network.	✓	✓	x	Medium
Keyboard and mouse sensor ⁵	Uses the desktop's PC touches and movements of keyboard and mouse to detect occupancy.	✓	x	x	n/a
Side-looking occupancy sensor ⁶	Increments or decrements the counter depending on entering and exiting movements in closed space.	✓	x	x	Low

A.2. Characteristics of occupancy detection sensors

Technology	Working principle	Functionality			Cost
		Detect	Count	Track	
Light ⁷	Marks indoor presence through light activation.	✓	x	x	n/a
Temperature ⁷	The difference between space's temperature with and without occupants.	✓	x	x	Low
Radar ⁸	Sends the radar waves to detect the moving object.	✓	x	x	Medium
Depth sensor ^{8,10}	Uses object-detection algorithms to identify occupants in a scene.	✓	✓	x	High
Infrared thermal camera ^{8,9}	Eliminates background through feature extraction from images.	✓	✓	x	High
Sound sensor ^{8,9}	Measures variation in sound waves as occupants interact with space.	✓	x	x	Low
GPS ^{8,9}	Monitors GPS signal as an indicator for occupant(s) absence from an indoor positioning environment.	✓	✓	✓	Medium
Door reed switch ^{8,9}	Magnetic switches detect a door's state of being closed or open.	✓	✓	x	Low
Ultra-Wideband Sensor (UWB) ⁹	Measures changes in amplitudes, frequencies and phases of multipath signals existing in radio channel.	✓	✓	✓	Medium

The cost criteria definition is as follows:

Low < 10 CHF, 10 CHF ≤ Medium ≤ 100 CHF, High > 100 CHF.

The n/a in costs signifies that such technology is assumed to be installed by default.

¹ [Chen 2018]

² [Salimi 2019]

³ [Yan 2015]

⁴ [Kleiminger 2015b]

⁵ [Zhao 2015]

⁶ [Weidman 2018]

⁷ [Candanedo 2016]

⁸ [Jung 2019]

⁹ [Priyadarshini 2015]

¹⁰ [Diraco 2015]

Table A.4 – Advantages and disadvantages of sensor technologies used for occupancy detection

Technology	Advantages	Disadvantages
PIR	<ul style="list-style-type: none"> • High accuracy 80%+ • Easy deployment • Low energy consumption 	<ul style="list-style-type: none"> • Misses static occupants • Requires setup, commissioning, calibration, and frequent maintenance • Occupants needs to be in sensor's proximity • Influenced by HVAC
Smart meter	<ul style="list-style-type: none"> • Non-intrusive for occupants • High accuracy 90%+ in residential and 80%+ in tertiary buildings • No extra cost if already present in the building 	<ul style="list-style-type: none"> • Data access needs to be provided • Requires extensive collection periods to identify occupancy patterns
BLE	<ul style="list-style-type: none"> • High accuracy 80%+ • Energy-efficient • Up to 100m broadcast range 	<ul style="list-style-type: none"> • Misses occupants not connected to Bluetooth
Chair sensor	<ul style="list-style-type: none"> • High accuracy 99% 	<ul style="list-style-type: none"> • Only suitable for spaces with chairs • Misses standing occupants
Vibration sensor	<ul style="list-style-type: none"> • High accuracy 87% • Distinguishes human-related signals 	<ul style="list-style-type: none"> • Requires large amount of sensors to adequately cover big spaces
Ultrasonic sensor	<ul style="list-style-type: none"> • High accuracy 95% 	<ul style="list-style-type: none"> • Triggered by false movements in adjacent spaces • Requires professional setup, commissioning, and tuning • Obscured by obstacles
UWB	<ul style="list-style-type: none"> • Device-free 	<ul style="list-style-type: none"> • Misses static occupants

A.2. Characteristics of occupancy detection sensors

Technology	Advantages	Disadvantages
RFID	<ul style="list-style-type: none"> • High accuracy up to 90% • Large communication range • Works without line of sight 	<ul style="list-style-type: none"> • Intrusive for occupants • Misses occupants without RFID tag • Prone to signal interference in small spaces
CO ₂ sensor	<ul style="list-style-type: none"> • Non-intrusive 	<ul style="list-style-type: none"> • Slow reaction time • Prone to influence of non-human nature, such as HVAC functioning, doors/windows opening • Works best in closed spaces • Requires calibration
Video camera	<ul style="list-style-type: none"> • High accuracy 96% 	<ul style="list-style-type: none"> • Intrusive for occupants • Requires storage, high bandwidth to transmit data, and high computational complexity to process it • Influenced by illumination • Requires clear line of sight • Limited observation angle
WiFi	<ul style="list-style-type: none"> • High accuracy 80% • Large communication range • Works without line of sight 	<ul style="list-style-type: none"> • Misses occupants not connected • Double counts occupants with multiple connected devices • Misinterprets static connected machines for occupants
Keyboard and mouse sensor	<ul style="list-style-type: none"> • High accuracy 92.8% • Easy to collect information through wireless networks • Enables data storage at PC itself 	<ul style="list-style-type: none"> • Misses occupants not using PCs • Privacy-concerning
Side-looking occupancy sensor	<ul style="list-style-type: none"> • Easy deployment 	<ul style="list-style-type: none"> • Misses occupants of small height • Occasionally mistakes leg movements for multiple occupants

Appendix A. Appendix

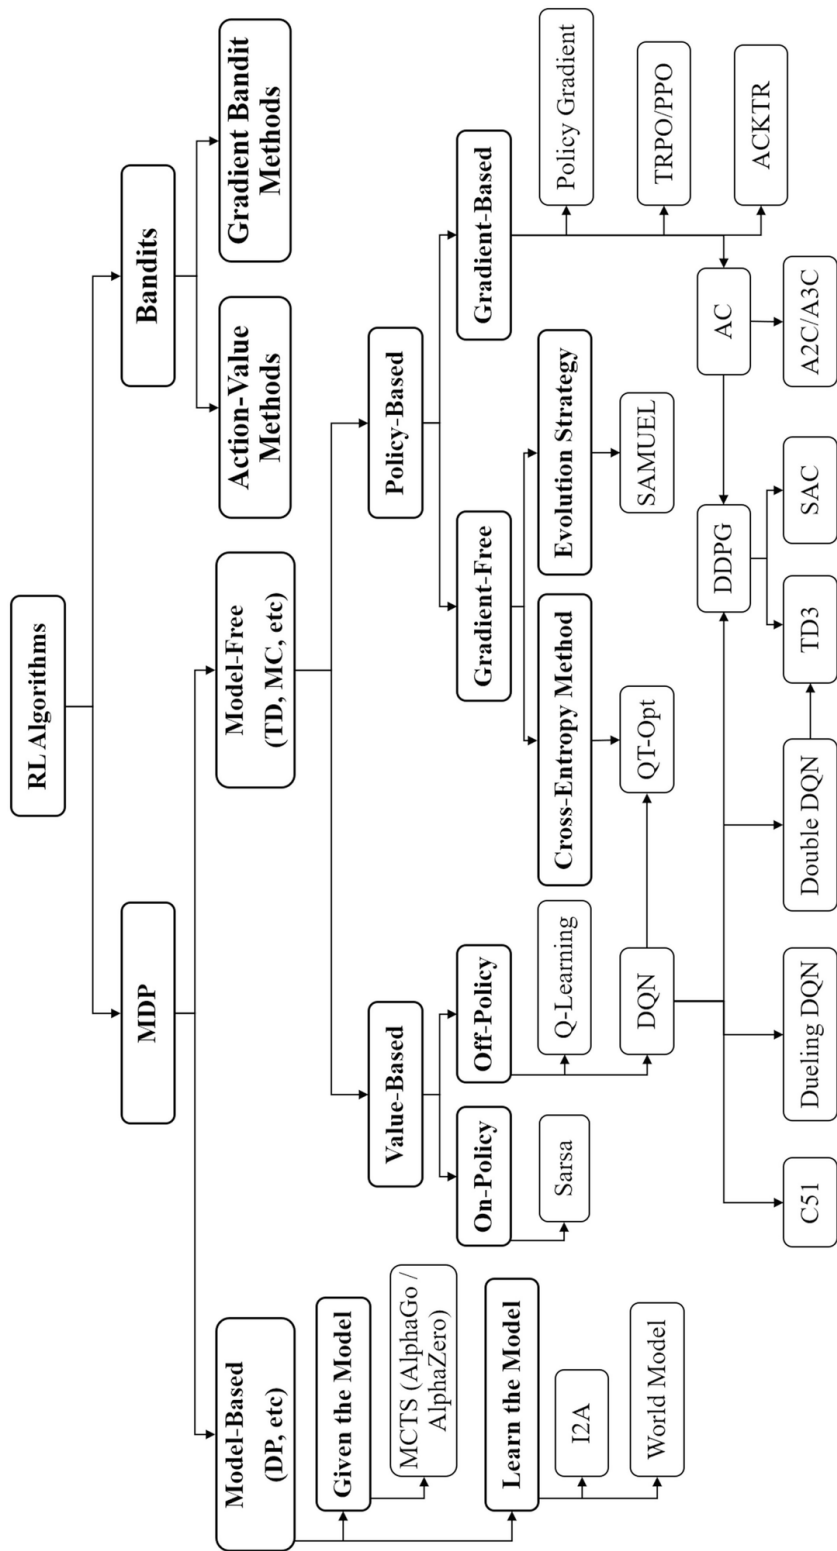
Technology	Advantages	Disadvantages
Light	<ul style="list-style-type: none">• High accuracy 80%	<ul style="list-style-type: none">• Not suitable for daytime• Responds to non-human triggers, such as daylight
Temperature	<ul style="list-style-type: none">• Accuracy 70%+	<ul style="list-style-type: none">• Prone to time delays• Responses to natural outdoor temperature cycling
Radar	<ul style="list-style-type: none">• Long-range distance• 360° view angle	<ul style="list-style-type: none">• Detects only moving objects• Wall penetration misjudges occupancy• Interferes with WiFi and Bluetooth
Infrared thermal camera	<ul style="list-style-type: none">• Requires no external source of light to maintain vision	<ul style="list-style-type: none">• Requires clear line of sight without obstructions• Requires high bandwidth to transmit data
Sound sensor	<ul style="list-style-type: none">• Easy to manipulate in real-time	<ul style="list-style-type: none">• Responds to non-human triggers• Interferes with radio signals
GPS	<ul style="list-style-type: none">• Provides direct localization	<ul style="list-style-type: none">• Requires physical device for precise indoor localization
Door reed switch	<ul style="list-style-type: none">• Low operational power• Easy to integrate	<ul style="list-style-type: none">• Useless if doors kept constantly open or closed
Depth sensor	<ul style="list-style-type: none">• Robust to light variation• Non-intrusive	<ul style="list-style-type: none">• Requires sufficient illumination• Requires specific computer vision techniques for data processing

A.3 Feature engineering in occupancy forecasting

Table A.5 – Sets of features used to transform raw data in supervised occupancy forecasting

Statistical	Load curve shape	Time-related
Power rounded	Load curve similarity	Day of the week
Quartiles	Area under curve	Weekday or weekend
Daily accumulated mean	Is peak	Month
Daily change to accumulated mean	Is valley	Season
First order difference	Discrete power level	Time of the day
Second order difference	Change to night mean	Float hour
Daily sum	Change to daily mean	Fourier transform
Hourly sum	Mean crossing	Holiday
Daily mean	Slope change	
Hourly mean		
Daily max		
Hourly max		
Daily min		
Hourly min		
Mean-to-max daily ratio		
Mean-to-max hourly ratio		
Min-to-mean daily ratio		
Min-to-mean hourly ratio		
Standard deviation daily		
Standard deviation hourly		
Variance daily		
Variance hourly		
Range daily		
Range hourly		
Median daily		
Median hourly		
Cumulative sum		
Squares		
Root		
Daily auto-correlation lag1		
Hourly auto-correlation lag1		
Daily raw-to-max ratio		
Hourly raw-to-max ratio		
Daily raw-to-mean ratio		
Hourly raw-to-mean ratio		
Daily change to standard deviation		
Cumulative sum of abs difference		
First order shifted difference		
Daily normalized power		
Daily max abs difference		

A.4 Taxonomy of reinforcement learning



Bibliography

- [Abnett 2020] Kate Abnett. *EU to target 30 million electric cars by 2030*. Reuters, 2020. <https://www.reuters.com/article/us-climate-change-eu-transport/eu-to-target-30-million-electric-cars-by-2030-draft-idINKBN28E2KM> [Accessed March 02, 2021].
- [Abousleiman 2014] R. Abousleiman and O. Rawashdeh. *Energy-efficient routing for electric vehicles using metaheuristic optimization frameworks*. In MELECON 2014 - 17th IEEE Mediterranean Electrotechnical Conference, pages 298–304. IEEE, 2014.
- [ACEA 2020] European Automobile Manufacturers Association ACEA. *New passenger car registrations by fuel type in the European Union: Quarter 4 2020*, 2020. https://www.acea.be/uploads/press_releases_files/20210204_PRPC_fuel_Q4_2020_FINAL.pdf [Accessed March 02, 2021].
- [Akbar 2015] Adnan Akbar, Michele Nati, Francois Carrez and Klaus Moessner. *Contextual occupancy detection for smart office by pattern recognition of electricity consumption data*. In Proceedings of the IEEE International Conference on Communications (ICC) 2015, pages 561–566, London, UK, 2015.
- [Alhamoud 2014] Alaa Alhamoud, Felix Ruettiger, Andreas Reinhardt, Frank Englert, Daniel Burgstahler, Doreen Bohnstedt, Christian Gottron and Ralf Steinmetz. *SMARTEN-ERGY.KOM: An intelligent system for energy saving in smart home*. In 39th Annual IEEE Conference on Local Computer Networks Workshops, pages 685–692. IEEE, September 2014.
- [Andoni 2019] Merlinda Andoni, Valentin Robu, David Flynn, Simone Abram, Dale Geach, David Jenkins, Peter McCallum and Andrew Peacock. *Blockchain technology in the energy sector: A systematic review of challenges and opportunities*. Renewable and Sustainable Energy Reviews, vol. 100, pages 143–174, 2019.
- [Androulaki 2018] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevichet *al.* *Hyperledger fabric: a distributed operating system for permissioned blockchains*. In Proceedings of the thirteenth EuroSys conference, pages 1–15, 2018.

Bibliography

- [Andrychowicz 2017] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel and Wojciech Zaremba. *Hindsight experience replay*. In Advances in neural information processing systems, pages 5048–5058, 2017.
- [Attwood 2017] J. Attwood, C. Curry and M. Wilshire. *Digitalization of Energy Systems*. Bloomberg New Energy Finance, pages 1–57, 2017.
- [Bailey 2020] Tucker Bailey, Adam Maruyama and Daniel Wallance. *The energy-sector threat: How to address cybersecurity vulnerabilities*. , 2020.
- [Bao 2020] Jiabin Bao, Debiao He, Min Luo and Kim-Kwang Raymond Choo. *A Survey of Blockchain Applications in the Energy Sector*. IEEE Systems Journal, pages 1–12, 2020.
- [Barter 2013] Paul Barter. *Cars are parked 95% of the time. Let's check!* Reinventing Parking, 2013. <https://www.reinventingparking.org/2013/02/cars-are-parked-95-of-time-lets-check.html> [Accessed March 8, 2021].
- [Batra 2016] Nipun Batra, Amarjeet Singh and Kamin Whitehouse. *Gemello: Creating a Detailed Energy Breakdown from Just the Monthly Electricity Bill*. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16, pages 431–440, New York, New York, USA, 2016. ACM Press.
- [Baza 2019] Mohamed Baza, Mahmoud Nabil, Muhammad Ismail, Mohamed Mahmoud, Erchin Serpedin and Mohammad Ashiqur Rahman. *Blockchain-based charging coordination mechanism for smart grid energy storage units*. In 2019 IEEE International Conference on Blockchain (Blockchain), pages 504–509. IEEE, 2019.
- [Beckel 2014] Christian Beckel, Wilhelm Kleiminger, Romano Cicchetti, Thorsten Staaake and Silvia Santini. *The ECO data set and the performance of non-intrusive load monitoring algorithms*. In Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings - BuildSys '14, pages 80–89, 2014.
- [Becker 2018] Vincent Becker and Wilhelm Kleiminger. *Exploring zero-training algorithms for occupancy detection based on smart meter measurements*. Computer Science - Research and Development, vol. 33, pages 25–36, 2018.
- [Bester 2019] Craig J. Bester, Steven D. James and George D. Konidaris. *Multi-Pass Q-Networks for Deep Reinforcement Learning with Parameterised Action Spaces*. arXiv preprint arXiv:1905.04388, 2019.
- [Biansoongnern 2016] Somchai Biansoongnern and Boonyang Plangklang. *Nonintrusive load monitoring (NILM) using an Artificial Neural Network in embedded system with low sampling rate*. In 2016 13th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology, ECTI-CON 2016, pages 1–4, June 2016.

- [Bonfigli 2015] Roberto Bonfigli, Stefano Squartini, Marco Fagiani and Francesco Piazza. *Unsupervised algorithms for non-intrusive load monitoring: An up-to-date overview*. In 2015 IEEE 15th International Conference on Environment and Electrical Engineering (EEEIC), pages 1175–1180. IEEE, June 2015.
- [Bozorgi 2017] Amir Masoud Bozorgi, Mehdi Farasat and Anas Mahmoud. *A time and energy efficient routing algorithm for electric vehicles based on historical driving data*. IEEE Transactions on Intelligent Vehicles, vol. 2, no. 4, pages 308–320, 2017.
- [Brockman 2016] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang and Wojciech Zaremba. *Openai gym*. arXiv preprint arXiv:1606.01540, 2016.
- [Bughin 2017] Jacques Bughin, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlstrom, Nicolaus Henke and Monica Trench. *Artificial Intelligence: The next digital frontier?*, pages 1–80, 2017.
- [Bughin 2018] Jacques Bughin, Jeongmin Seong, James Manyika, Michael Chui and Raoul Joshi. *Notes from the AI frontier: modeling the impact of AI on the world economy*. McKinsey Global Institute, pages 1–64, 2018.
- [Busoniu 2010] L. Busoniu, R. Babuska and B. De Schutter. *Multi-agent reinforcement learning: An overview*. Studies in Computational Intelligence, vol. 310, pages 183–221, 2010.
- [Buterin 2014] Vitalik Buterin et al. *A next-generation smart contract and decentralized application platform*. White paper, vol. 3, no. 37, 2014.
- [Candanedo 2016] Luis M Candanedo and Véronique Feldheim. *Accurate occupancy detection of an office room from light, temperature, humidity and CO2 measurements using statistical learning models*. Energy and Buildings, vol. 112, pages 28–39, 2016.
- [Candanedo 2017] Luis M. Candanedo, Véronique Feldheim and Dominique Deramaix. *A methodology based on Hidden Markov Models for occupancy detection and a case study in a low energy residential building*. Energy and Buildings, vol. 148, pages 327–341, 2017.
- [Cauwer 2019] Cedric De Cauwer, Wouter Verbeke, Joeri Van Mierlo and Thierry Coosemans. *A model for range estimation and energy-efficient routing of electric vehicles in real-world conditions*. IEEE Transactions on Intelligent Transportation Systems, vol. 21, pages 1–14, 2019.
- [Chen 2013] Dong Chen, Sean Barker, Adarsh Subbaswamy, David Irwin and Prashant Shenoy. *Non-intrusive occupancy monitoring using smart meters*. In Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings - BuildSys’13, pages 1–8, Rome, Italy, 2013.

- [Chen 2018] Zhenghua Chen, Chaoyang Jiang and Lihua Xie. *Building occupancy estimation and detection: A review*. Energy and Buildings, vol. 169, pages 260–270, 2018.
- [Chen 2020a] Huashan Chen, Marcus Pendleton, Laurent Njilla and Shouhuai Xu. *A survey on ethereum systems security: Vulnerabilities, attacks, and defenses*. ACM Computing Surveys (CSUR), vol. 53, no. 3, pages 1–43, 2020.
- [Chen 2020b] Xi Chen, Tianyang Zhang, Wenxing Ye, Zhiwei Wang and Herbert Ho-Ching Iu. *Blockchain-based Electric Vehicle Incentive System for Renewable Energy Consumption*. IEEE Transactions on Circuits and Systems II: Express Briefs, 2020.
- [Cheng 2019] Lefeng Cheng and Tao Yu. *A new generation of AI: A review and perspective on machine learning technologies applied to smart energy and electric power systems*. International Journal of Energy Research, vol. 43, pages 1928–1973, 2019.
- [Chiş 2016] Adriana Chiş, Jarmo Lundén and Visa Koivunen. *Reinforcement learning-based plug-in electric vehicle charging with forecasted price*. IEEE Transactions on Vehicular Technology, vol. 66, no. 5, pages 3674–3684, 2016.
- [Claessens 2013] Bert J Claessens, Stijn Vandael, Frederik Ruelens, Klaas De Craemer and Bart Beusen. *Peak shaving of a heterogeneous cluster of residential flexibility carriers using reinforcement learning*. In IEEE PES ISGT Europe 2013, pages 1–5, Copenhagen, Denmark, 2013.
- [Clement 2020] J. Clement. *Share of global mobile website traffic 2015-2020*. Statista: Mobile internet usage worldwide, 2020. <https://www.statista.com/statistics/277125/share-of-website-traffic-coming-from-mobile-devices> [Accessed January 11, 2021].
- [Commission 2020] European Commission. *Heating and cooling*. Energy Efficiency, 2020. <https://ec.europa.eu/energy/topics/energy-efficiency/heating-and-cooling> [Accessed June 3, 2020].
- [Daanish 2017] A. R. Daanish and B. K. Naick. *Implementation of charging station based electric vehicle routing problem using nearest neighbour search algorithm*. In 2017 2nd IEEE International Conference on Intelligent Transportation Engineering (ICITE), pages 52–56. IEEE, 2017.
- [Dai 2020] Xilei Dai, Junjie Liu and Xin Zhang. *A review of studies applying machine learning models to predict occupancy and window-opening behaviours in smart buildings*. Energy and Buildings, vol. 214, page 110159, 2020.
- [Dalen 2013] Hanne Marit Dalen and Bodil M. Larsen. *Residential end-use electricity demand: Development over time*. Statistics Norway, discussion paper, no. 736, pages 1–24, 2013.
- [Delzendeh 2017] Elham Delzendeh, Song Wu, Angela Lee and Ying Zhou. *The impact of occupants’ behaviours on building energy analysis: a research review*. Renewable and Sustainable Energy Reviews, vol. 80, pages 1061–1071, 2017.

- [Desjardins 2018] Jeff Desjardins. *How Long Does It Take to Hit 50 Million Users?* Visual Capitalist, 2018. <https://www.visualcapitalist.com/how-long-does-it-take-to-hit-50-million-users> [Accessed January 11, 2021].
- [Di Silvestre 2018a] Maria Luisa Di Silvestre, Salvatore Favuzza, Eleonora Riva Sanseverino and Gaetano Zizzo. *How Decarbonization, Digitalization and Decentralization are changing key power infrastructures*. Renewable and Sustainable Energy Reviews, vol. 93, pages 483–498, 2018.
- [Di Silvestre 2018b] Maria Luisa Di Silvestre, Pierluigi Gallo, Mariano Giuseppe Ippolito, Eleonora Riva Sanseverino and Gaetano Zizzo. *A technical approach to the energy blockchain in microgrids*. IEEE Transactions on Industrial Informatics, vol. 14, no. 11, pages 4792–4803, 2018.
- [Di Wu 2018] Rabusseau Guillaume Di Wu, Francois lavet Vincent, Precup Doina and Boulet Benoit. *Optimizing home energy management and electric vehicle charging with reinforcement learning*. In 2018 16th Workshop at the Federated AI Meeting: Adaptive Learning Agents, pages 1–8, Stockholm, Sweden, 2018.
- [Digiconomist 2021] Digiconomist. *Ethereum Energy Consumption Index*, 2021. <https://digiconomist.net/ethereum-energy-consumption/> [Accessed March 24, 2021].
- [Dijkstra 1959] Edsger Wybe Dijkstra. *A note on two problems in connexion with graphs*. Numerische Mathematik, vol. 1, pages 269–271, 1959.
- [Ding 2020] Tao Ding, Ziyu Zeng, Jiawen Bai, Boyu Qin, Yongheng Yang and Mohammad Shahidehpour. *Optimal Electric Vehicle Charging Strategy with Markov Decision Process and Reinforcement Learning Technique*. IEEE Transactions on Industry Applications, vol. 9994, 2020.
- [Diraco 2015] Giovanni Diraco, Alessandro Leone and Pietro Siciliano. *People occupancy detection and profiling with 3D depth sensors for building energy management*. Energy and Buildings, vol. 92, pages 246–266, 2015.
- [Dong 2014] B. Dong and K. Poh Lam. *A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting*. Building Simulation, vol. 7, pages 89–106, 2014.
- [Dunkan 2019] Luke Dunkan, John Light, Luis Ivan Cuende, Jorge Izquierdo and Facu Spagnuolo. *Aragon Network White Paper*, 2019. <https://github.com/aragon/whitepaper> [Accessed October 20, 2020].
- [EC 2020] European Commission EC. *2030 Climate & Energy Framework*, 2020. https://ec.europa.eu/clima/policies/strategies/2030_en [Accessed March 10, 2021].

Bibliography

- [EC 2021] European Commission EC. *Energy performance of buildings directive*, 2021. https://ec.europa.eu/energy/topics/energy-efficiency/energy-efficient-buildings/energy-performance-buildings-directive_en [Accessed January 11, 2021].
- [EEA 2020a] European Environment Agency EEA. *Final energy consumption by sector and fuel in Europe*. European Environment Agency Indicators , 2020. <https://www.eea.europa.eu/data-and-maps/indicators/final-energy-consumption-by-sector-10/assessment> [Accessed June 3, 2020].
- [EEA 2020b] European Environmental Agency EEA. *Greenhouse gas emissions from transport in Europe*, 2020. <https://www.eea.europa.eu/data-and-maps/indicators/transport-emissions-of-greenhouse-gases-7/assessment> [Accessed March 10, 2021].
- [EEMA 2017] European Automobile Manufacturers Association EEMA. *Worldwide Harmonised Light Vehicle Test Procedure (WLTP)*, 2017. <https://www.wltpfacts.eu/what-is-wltp-how-will-it-work/> [Accessed April 27, 2020].
- [Elkhouchi 2019] Hamza Elkhouchi, Mohamed Bakhouya, Majdoulayne Hanifi and Driss El Ouadghiri. *On the use of Deep Learning Approaches for Occupancy prediction in Energy Efficient Buildings*. In Proceedings of 7th International Renewable and Sustainable Energy Conference, IRSEC 2019, pages 1–6, Agadir, Morocco, 2019.
- [Energy Web Foundation 2019] Energy Web Foundation. *The Energy Web Chain: Accelerating the Energy Transition with an Open-source, Decentralized Blockchain Platform*, 2019. <https://www.energyweb.org/reports/the-energy-web-chain> [Accessed February 02, 2021].
- [Energy Web 2020] Energy Web. *EW-DOS: The Energy Web Decentralized Operating System. An Open-Source Technology Stack to Accelerate the Energy Transition. Part II, Technology detail*, 2020. <https://www.energyweb.org/reports/EWDOS-Technology-Detail> [Accessed February 02, 2021].
- [Esa 2016] Nur Farahin Esa, Md Pauzi Abdullah and Mohammad Yusri Hassan. *A review disaggregation method in Non-intrusive Appliance Load Monitoring*. Renewable and Sustainable Energy Reviews, vol. 66, pages 163–173, 2016.
- [Ethereum 2021] Ethereum. *Ethereum Gas Station*, 2021. <https://ethgasstation.info/> [Accessed February 02, 2021].
- [Etherscan 2021] Etherscan. *Ropsten Testnet Explorer*, 2021. <https://ropsten.etherscan.io> [Accessed February 02, 2021].
- [EU 2020] European Union EU. *Stepping up Europe's 2030 climate ambition*. Communication from the commission to the European parliament, the council, the European economic and social committee and the committee of the regions, pages 1–26, 2020.

- [Eurostat 2021] Eurostat. *Electricity prices for non-household consumers - bi-annual data (from 2007 onwards)*. Eurostat Data Browser, 2021. https://ec.europa.eu/eurostat/databrowser/view/nrg_pc_205/default/bar?lang=en [Accessed February 25, 2021].
- [EV Pass 2021] EV Pass. *The EVPass subscription fees*, 2021. <https://www.evpass.ch/Subscription> [Accessed March 02, 2021].
- [EVD 2021a] Electric Vehicle Database EVD. *Citroen C-Zero*, 2021. <https://ev-database.org/car/1094/Citroen-C-Zero> [Accessed March 15, 2021].
- [EVD 2021b] Electric Vehicle Database EVD. *Electric Vehicle Database*, 2021. <https://ev-database.org/> [Accessed March 02, 2021].
- [EVD 2021c] Electric Vehicle Database EVD. *Hyundai Kona Electric 64 kWh*, 2021. <https://ev-database.org/car/1204/Hyundai-Kona-Electric-64-kWh> [Accessed March 02, 2021].
- [EWF 2021] Energy Web Foundation EWF. *Energy Web Chain Explorer*, 2021. <https://explorer.energyweb.org/txs> [Accessed April 27, 2021].
- [Faheem 2018] M. Faheem, S. B.H. Shah, R. A. Butt, B. Raza, M. Anwar, M. W. Ashraf, Md A. Ngadi and V. C. Gungor. *Smart grid communication and information technologies in the perspective of Industry 4.0: Opportunities and challenges*. Computer Science Review, vol. 30, pages 1–30, 2018.
- [Fan 2019] Zhou Fan, Rui Su, Weinan Zhang and Yong Yu. *Hybrid actor-critic reinforcement learning in parameterized action space*. In IJCAI International Joint Conference on Artificial Intelligence, pages 2279–2285, Macao, China, 2019.
- [Fang 2019] Xiaohan Fang, Jinkuan Wang, Guanru Song, Yinghua Han, Qiang Zhao and Zhiao Cao. *Multi-agent reinforcement learning approach for residential microgrid energy scheduling*. Energies, vol. 13, pages 1–26, 2019.
- [Faustine 2017] Anthony Faustine, Nerey Henry Mvungi, Shubi Kaijage and Kisangiri Michael. *A Survey on Non-Intrusive Load Monitoring Methodies and Techniques for Energy Dis-aggregation Problem*. ArXiv, March 2017.
- [Felipe 2014] Angel Felipe, M. Teresa Ortuño, Giovanni Righini and Gregorio Tiradoa. *A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges*. Transportation Research Part E: Logistics and Transportation Review, vol. 71, no. 0, pages 111–128, 2014.
- [Feng 2020] Cong Feng, Ali Mehmani and Jie Zhang. *Deep Learning-based Real-time Building Occupancy Detection Using AMI Data*. IEEE Transactions on Smart Grid, vol. 3053, pages 1–12, 2020.

Bibliography

- [Fu 2020] Zhengtang Fu, Peiwu Dong and Yanbing Ju. *An intelligent electric vehicle charging system for new energy companies based on consortium blockchain*. Journal of Cleaner Production, page 121219, 2020.
- [Gellings 2017] W. Clark Gellings. *Evolving practice of demand-side management*. Journal of Modern Power Systems and Clean Energy, no. 5, pages 1–9, 2017.
- [Girardin 2012] L. Girardin. *A GIS-based methodology for the evaluation of integrated energy systems in urban area*. PhD thesis, EPFL, 2012.
- [Gorenflo 2019] Christian Gorenflo, Stephen Lee, Lukasz Golab and Srinivasan Keshav. *Fast-fabric: Scaling hyperledger fabric to 20,000 transactions per second*. In 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), pages 455–463. IEEE, 2019.
- [Grijpink 2020] Ferry Grijpink, Eric Kutcher, Alexandre Menard, Sree Ramaswamy, Davide Schiavotto, James Manyika, Michael Chui, Rob Hamill and Emir Okan. *Connected world: An evolution in connectivity beyond the 5G revolution*. , pages 1–100, 2020.
- [Guan 2015] Chenxiao Guan, Yanzhi Wang, Xue Lin, Shahin Nazarian and Massoud Pedram. *Reinforcement learning-based control of residential energy storage systems for electric bill minimization*. In 2015 12th Annual IEEE Consumer Communications and Networking Conference (CCNC), pages 637–642, Las Vegas, USA, 2015. IEEE.
- [Habib 2017] Usman Habib and Gerhard Zucker. *Automatic occupancy prediction using unsupervised learning in buildings data*. In Proceedings of the IEEE International Symposium on Industrial Electronics (ISIE 2017), pages 1471–1476, Edinburgh, Scotland, 2017.
- [Habibi 2013] Mozhi Habibi. *Cloud computing for energy and utilities: Driving IT transformation and smarter energy systems*. IBM: Energy & Utilities, pages 1–12, 2013.
- [Hall 2018] Dale Hall and Nic Lutsey. *Effects of battery manufacturing on electric vehicle life-cycle greenhouse gas emissions*, 2018. <https://theicct.org/publications/EV-battery-manufacturing-emissions> [Accessed March 18, 2021].
- [Hart 1992] George William Hart. *Nonintrusive appliance load monitoring*. Proceedings of the IEEE, vol. 80, no. 12, pages 1870–1891, 1992.
- [Hassan 2019] Naveed Ul Hassan, Chau Yuen and Dusit Niyato. *Blockchain technologies for smart energy systems: Fundamentals, challenges, and solutions*. IEEE Industrial Electronics Magazine, vol. 13, no. 4, pages 106–118, 2019.
- [Hasselt 2010] Hado V Hasselt. *Double Q-learning*. In Advances in neural information processing systems 23 (NIPS), pages 2613–2621, Vancouver, Canada, 2010.
- [Hausknecht 2015] Matthew Hausknecht and Peter Stone. *Deep reinforcement learning in parameterized action space*. arXiv preprint arXiv:1511.04143, 2015.

- [He 2016] Kanghang He, Lina Stankovic, Jing Liao and Vladimir Stankovic. *Non-Intrusive Load Disaggregation using Graph Signal Processing*. IEEE Transactions on Smart Grid, vol. 9, no. 3, pages 1739–1747, 2016.
- [Henderson 2017] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup and David Meger. *Deep reinforcement learning that matters*. arXiv preprint arXiv:1709.06560, 2017.
- [Henze 2020] Veronika Henze. *Battery Pack Prices Cited Below \$100/kWh for the First Time in 2020, while Market Average sits at \$137/kWh*. Bloomberg New Energy Finance, 2020. <https://about.bnef.com/blog/battery-pack-prices-cited-below-100-kwh-for-the-first-time-in-2020-while-market-average-sits-at-137-kwh> [Accessed January 11, 2021].
- [Henzelmann 2018] Torster Henzelmann, Frederik Hammermeister, Sebastian Preiss, Bernhard Wurm, Karsten Schroer and Lukas Nonnenmacher. *Artificial intelligence: A smart move for utilities*. Roland Berger, pages 1–19, 2018.
- [Hofman 2010] T Hofman and CH Dai. *Energy efficiency analysis and comparison of transmission technologies for an electric vehicle*. In 2010 IEEE vehicle power and propulsion conference, pages 1–6. IEEE, 2010.
- [Holweger 2019] Jordan Holweger, Marina Dorokhova, Lionel Bloch, Christophe Ballif and Nicolas Wyrsh. *Unsupervised algorithm for disaggregating low-sampling-rate electricity consumption of households*. Sustainable Energy, Grids and Networks, vol. 19, page 100244, 2019.
- [Hoover 2021] Zealan Hoover, Florian Nägele, Evan Polymeneas and Shivika Sahdev. *How charging in buildings can power up the electric-vehicle industry*. McKinsey, 2021. <https://www.mckinsey.com/industries/electric-power-and-natural-gas/our-insights/how-charging-in-buildings-can-power-up-the-electric-vehicle-industry> [Accessed March 18, 2021].
- [Howard 2019] Bianca Howard, Salvador Acha, Nilay Shah and John Polak. *Implicit Sensing of Building Occupancy Count with Information and Communication Technology Data Sets*. Building and Environment, vol. 157, pages 297–308, 2019.
- [Huang 2018] Xiaohong Huang, Cheng Xu, Pengfei Wang and Hongzhe Liu. *LNSC: A security model for electric vehicle and charging pile management based on blockchain ecosystem*. IEEE Access, vol. 6, pages 13565–13574, 2018.
- [Huang 2019] Xiaohong Huang, Yong Zhang, Dandan Li and Lu Han. *An optimal scheduling algorithm for hybrid EV charging scenario using consortium blockchains*. Future Generation Computer Systems, vol. 91, pages 555–562, 2019.
- [Huchuk 2019] Brent Huchuk, Scott Sanner and William O’Brien. *Comparison of machine learning models for occupancy prediction in residential buildings using connected thermostat data*. Building and Environment, vol. 160, page 106177, 2019.

Bibliography

- [Hyperledger 2021] Hyperledger. *Hyperledger Architecture, Volume 1: Introduction to Hyperledger business blockchain design philosophy and consensus*, 2021. https://www.hyperledger.org/wp-content/uploads/2017/08/Hyperledger_Arch_WG_Paper_1_Consensus.pdf [Accessed February 02, 2021].
- [IEA 2017] International Energy Agency IEA. *Digitalization & Energy*, 2017.
- [IEA 2020a] International Energy Agency IEA. *Global EV Outlook*, 2020. <https://www.iea.org/reports/global-ev-outlook-2020> [Accessed July 21, 2020].
- [IEA 2020b] International Energy Agency IEA. *World Energy Outlook 2020*, 2020.
- [Innosuisse 2019] Innosuisse. *Digitalisation in energy and mobility via SC-CER*, 2019. https://www.innosuisse.ch/inno/en/home/promotion-initiatives/impulsprogramm_digitalisierung.html [Accessed December 28, 2020].
- [IRENA 2020] International Renewable Energy Agency IRENA. *Renewable power generation costs in 2019*, 2020.
- [Izquierdo 2019] Jorge Izquierdo. *Aragon Chain: a proof of stake blockchain for the Aragon community*, 2019. <https://blog.aragon.one/aragon-chain> [Accessed October 20, 2020].
- [Jindal 2018] A. Jindal, N. Kumar and J. Rodrigues. *A heuristic-based smart HVAC energy management scheme for university buildings*. IEEE Transactions on Industrial Informatics, vol. 3203, pages 1–12, 2018.
- [Jung 2019] Wooyoung Jung and Farrokh Jazizadeh. *Human-in-the-loop HVAC operations: A quantitative review on occupancy, comfort, and energy-efficiency dimensions*. Applied Energy, vol. 239, pages 1471–1508, 2019.
- [Kang 2017] Jiawen Kang, Rong Yu, Xumin Huang, Sabita Maharjan, Yan Zhang and Ekram Hossain. *Enabling localized peer-to-peer electricity trading among plug-in hybrid electric vehicles using consortium blockchains*. IEEE Transactions on Industrial Informatics, vol. 13, no. 6, pages 3154–3164, 2017.
- [Kathirgamanathan 2021] Anjukan Kathirgamanathan, Mattia De Rosa, Eleni Mangina and Donal P. Finn. *Data-driven predictive control for unlocking building energy flexibility: A review*. Renewable and Sustainable Energy Reviews, vol. 135, page 110120, 2021.
- [Katrenko 2020] Anna Katrenko and Mihail Sotnichek. *Blockchain Attack Vectors: Vulnerabilities of the Most Secure Technology*, 2020. <https://www.apriorit.com/dev-blog/578-blockchain-attack-vectors> [Accessed October 20, 2020].
- [Kelly 2015a] Jack Kelly and William Knottenbelt. *Neural NILM: Deep neural networks applied to energy disaggregation*. In Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments, pages 55–64. ACM, 2015.

- [Kelly 2015b] Jack Kelly and William Knottenbelt. *The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes*. Scientific Data, vol. 2, page 150007, March 2015.
- [Kelly 2018] Jack Kelly. *Appliance power demand specifications*. NILM Wiki, 2018. http://wiki.nilme.eu/index.php?title=Appliance_power_demand [Accessed January 11, 2021].
- [Kim 2011] Hyungsul Kim, Manish Marwah, Martin Arlitt, Geoff Lyon and Jiawei Han. *Un-supervised Disaggregation of Low Frequency Power Measurements*. In Proceedings of the 2011 SIAM International Conference on Data Mining, pages 747–758. Society for Industrial and Applied Mathematics, Philadelphia, PA, April 2011.
- [Kim 2018] Sunyong Kim and Hyuk Lim. *Reinforcement learning based energy management algorithm for smart energy buildings*. Energies, vol. 11, no. 8, page 2010, 2018.
- [Kleiminger 2014] W. Kleiminger, F. Mattern and S. Santini. *Predicting household occupancy for smart heating control: A comparative performance analysis of state-of-the-art approaches*. Energy and Buildings, vol. 85, pages 493–505, 2014.
- [Kleiminger 2015a] Wilhelm Kleiminger. *Occupancy sensing and prediction for automated energy savings*. PhD thesis, ETH Zurich, 2015.
- [Kleiminger 2015b] Wilhelm Kleiminger, Christian Beckel and Silvia Santini. *Household occupancy monitoring using electricity meters*. In Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp’15, pages 975–986, Osaka, Japan, 2015. ACM.
- [Klepeis 2001] N. E. Klepeis, W. C. Nelson, W. R. Ott, J. P. Robinson and A.M Tsang. *The National Human Activity Pattern Survey (NHAPS): a resource for assessing exposure to environmental pollutants*. Journal of Exposure Science & Environmental Epidemiology, vol. 11, pages 231–252, 2001.
- [Kloppenborg 2019] Sanneke Kloppenborg and Marten Boekelo. *Digital platforms and the future of energy provisioning: Promises and perils for the next phase of the energy transition*. Energy Research & Social Science, vol. 49, pages 68–73, 2019.
- [Knirsch 2018] Fabian Knirsch, Andreas Unterweger and Dominik Engel. *Privacy-preserving blockchain-based electric vehicle charging with dynamic tariff decisions*. Computer Science-Research and Development, vol. 33, no. 1-2, pages 71–79, 2018.
- [Kolter 2010] J Zico Kolter, Siddharth Batra and Andrew Y. Ng. *Energy disaggregation via discriminative sparse coding*. In Advances in Neural Information Processing Systems, pages 1153–1161, 2010.
- [Krug 2021] Alexander Krug, Thomas Knoblinger and Florian Saeftel. *Electric vehicle charging in Europe*. Arthur D. Little Switzerland, 2021.

- <https://www.adlittle.ch/en/insights/viewpoints/electric-vehicle-charging-europe>
[Accessed March 02, 2021].
- [Kumar 2016] Kriti Kumar, Rahul Sinha, M. Girish Chandra and Naveen Kumar Thokala. *Data-driven electrical load disaggregation using graph signal processing*. In 2016 IEEE Annual India Conference (INDICON), pages 1–6. IEEE, December 2016.
- [Kumar 2017] Kriti Kumar and M. Girish Chandra. *An intuitive explanation of graph signal processing-based electrical load disaggregation*. In 2017 IEEE 13th International Colloquium on Signal Processing & its Applications (CSPA), pages 100–105. IEEE, March 2017.
- [Lange 2020] Steffen Lange, Johanna Pohl and Tilman Santarius. *Digitalization and energy consumption. Does ICT reduce energy demand?* Ecological Economics, vol. 176, page 106760, 2020.
- [Lasla 2020] Noureddine Lasla, Maryam Al-Ammari, Mohamed Abdallah and Mohamed Younis. *Blockchain based trading platform for electric vehicle charging in smart cities*. IEEE Open Journal of Intelligent Transportation Systems, vol. 1, pages 80–92, 2020.
- [Lee 2019] Sangyoon Lee and Dae-Hyun Choi. *Reinforcement learning-based energy management of smart home with rooftop solar photovoltaic system, energy storage system, and home appliances*. Sensors, vol. 19, no. 18, page 3937, 2019.
- [Lee 2020] Jaehyun Lee, Eunjung Lee and Jinho Kim. *Electric Vehicle Charging and Discharging Algorithm Based on Reinforcement Learning with Data-Driven Approach in Dynamic Pricing Scheme*. Energies, vol. 13, no. 8, page 1950, 2020.
- [Leising 2017] Matthew Leising. *The Ether thief*. Bloomberg, 2017. <https://www.bloomberg.com/features/2017-the-ether-thief> [Accessed October 20, 2020].
- [Li 2016] Wen Li, Patric Stanula, Patricia Egede, Sami Kara and Christoph Herrmann. *Determining the main factors influencing the energy consumption of electric vehicles in the usage phase*. Procedia CIRP, vol. 48, pages 352–357, 2016.
- [Li 2019] Yuancheng Li and Baiji Hu. *An Iterative Two-Layer Optimization Charging and Discharging Trading Scheme for Electric Vehicle Using Consortium Blockchain*. IEEE Transactions on Smart Grid, vol. 11, no. 3, pages 2627–2637, 2019.
- [Li 2020] Hepeng Li, Zhiqiang Wan and Haibo He. *A Deep Reinforcement Learning Based Approach for Home Energy Management System*. In 2020 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), pages 1–5, Washington, USA, 2020. IEEE.
- [Liang 2016] X. Liang, T. Hong and G. Shen. *Occupancy data analytics and prediction: a case study*. Building and Environment, vol. 102, pages 179–192, 2016.

- [Liao 2014] Jing Liao, Georgia Elafoudi, Lina Stankovic and Vladimir Stankovic. *Non-intrusive appliance load monitoring using low-resolution smart meter data*. In 2014 IEEE International Conference on Smart Grid Communications (SmartGridComm), pages 535–540. IEEE, November 2014.
- [Liessner 2019] Roman Liessner, Jakob Schmitt, Ansgar Dietermann and Bernard Bäker. *Hyperparameter Optimization for Deep Reinforcement Learning in Vehicle Energy Management*. In ICAART 2019 - 11th International Conference on Agents and Artificial Intelligence, pages 134–144, Prague, Czech Republic, 2019.
- [Lillicrap 2015] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver and Daan Wierstra. *Continuous control with deep reinforcement learning*. arXiv preprint arXiv:1509.02971, 2015.
- [Liu 2018] Chao Liu, Kok Keong Chai, Xiaoshuai Zhang, Eng Tseng Lau and Yue Chen. *Adaptive blockchain-based electric vehicle participation scheme in smart grid platform*. IEEE Access, vol. 6, pages 25657–25665, 2018.
- [Luthander 2015] Rasmus Luthander, Joakim Widén, Daniel Nilsson and Jenny Palm. *Photovoltaic self-consumption in buildings: A review*. Applied energy, vol. 142, pages 80–94, 2015.
- [Madina 2016] Carlos Madina, Inmaculada Zamora and Eduardo Zabala. *Methodology for assessing electric vehicle charging infrastructure business models*. Energy Policy, vol. 89, pages 284–293, 2016.
- [Makonin 2015] Stephen Makonin and Fred Popowich. *Nonintrusive load monitoring (NILM) performance evaluation*. Energy Efficiency, vol. 8, no. 4, pages 809–814, July 2015.
- [Mao 2020] Huiting Mao, Jianmai Shi, Yuzhen Zhou and Guoqing Zhang. *The electric vehicle routing problem with time windows and multiple recharging options*. IEEE Access, vol. 8, pages 114864–114875, 2020.
- [Marra 2012] Francesco Marra, Guang Ya Yang, Chresten Træholt, Esben Larsen, Claus Nygaard Rasmussen and Shi You. *Demand profile study of battery electric vehicle under different charging options*. In 2012 IEEE Power and Energy Society General Meeting, pages 1–7. IEEE, 2012.
- [Martinson 2021] Yann Martinson. *Smart charging of electric vehicles under uncertainty*, 2021. <https://infoscience.epfl.ch/record/283826?ln=en> [Accessed April 24, 2021].
- [Mason 2019] Karl Mason and Santiago Grijalva. *A review of reinforcement learning for autonomous building energy management*. Computers & Electrical Engineering, vol. 78, pages 300–312, 2019.
- [Masson 2015] Warwick Masson, Pravesh Ranchod and George Konidaris. *Reinforcement learning with parameterized actions*. arXiv preprint arXiv:1509.01644, 2015.

Bibliography

- [McKinsey 2018] McKinsey. *The Digital Utility: New challenges, capabilities, and opportunities*. McKinsey's Digital Utility Compendium, pages 1–76, 2018.
- [Miglani 2020] Arzoo Miglani, Neeraj Kumar, Vinay Chamola and Sherali Zeadally. *Blockchain for Internet of Energy management: Review, solutions, and challenges*. Computer Communications, vol. 151, pages 395–418, 2020.
- [Mirakhorli 2016] Amin Mirakhorli and Dong Bing. *Occupancy behavior based model predictive control for building indoor climate — a critical review*. Energy & Buildings, vol. 129, pages 499–513, 2016.
- [Mocanu 2018] Elena Mocanu, Decebal Constantin Mocanu, Phuong H Nguyen, Antonio Liotta, Michael E Webber, Madeleine Gibescu and Johannes G Slootweg. *On-line building energy optimization using deep reinforcement learning*. IEEE transactions on smart grid, vol. 10, no. 4, pages 3698–3708, 2018.
- [Morlock 2019] Florian Morlock, Bernhard Rolle, Michel Bauer and Oliver Sawodny. *Time optimal routing of electric vehicles under consideration of available charging infrastructure and a detailed consumption model*. IEEE Transactions on Intelligent Transportation Systems, vol. October, pages 1–13, 2019.
- [Mortier 2020] Thierry Mortier. *Why artificial intelligence is a game-changer for renewable energy*. 56th edition of the Renewable Energy Country Attractiveness Index (RECAI), 2020.
- [Nachum 2017] Ofir Nachum, Mohammad Norouzi, Kelvin Xu and Dale Schuurmans. *Bridging the gap between value and policy based reinforcement learning*. arXiv preprint arXiv:1702.08892, 2017.
- [Nair 2018] Ashvin Nair, Bob McGrew, Marcin Andrychowicz, Wojciech Zaremba and Pieter Abbeel. *Overcoming exploration in reinforcement learning with demonstrations*. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 6292–6299, Brisbane, Australia, 2018. IEEE.
- [Nakamoto 2008] Satoshi Nakamoto. *Bitcoin: A peer-to-peer electronic cash system*. White paper, 2008.
- [Nazari 2018] Mohammadreza Nazari, Afshin Oroojlooy, Martin Takác and Lawrence V. Snyder. *Reinforcement learning for solving the vehicle routing problem*. In Advances in Neural Information Processing Systems 31 (NIPS 2018), pages 9839–9849. NIPS Foundation, Inc., 2018.
- [NCTA 2020] The Internet & Television Association NCTA. *Growth In The Internet of Things*, 2020. <https://www.ncta.com/chart/growth-in-the-internet-of-things> [Accessed January 11, 2021].

- [Noel 2020] Lance Noel, Gerardo Zarazua de Rubens, Johannes Kester and Benjamin K. Sovacool. *Understanding the socio-technical nexus of Nordic electric vehicle (EV) barriers: A qualitative discussion of range, price, charging and knowledge*. Energy Policy, vol. 138, page 111292, 2020.
- [OFS 2015] Office fédéral de la statistique, section Mobilité OFS. *Population's transport behaviour*, 2015. <https://www.bfs.admin.ch/bfs/en/home/statistics/mobility-transport/passenger-transport/travel-behaviour.html> [Accessed March 02, 2021].
- [OFS 2020] Office fédéral de la statistique, section Mobilité OFS. *Véhicules routiers – parc, taux de motorisation*, 2020. <https://www.bfs.admin.ch/bfs/fr/home/statistiques/mobilite-transports/infrastructures-transport-vehicules/vehicules/vehicules-routiers-parc-taux-motorisation.html> [Accessed March 02, 2021].
- [Open Charge Map 2021] Open Charge Map. *Charging networks in Switzerland*, 2021. <https://openchargemap.org/site/country/switzerland/networks> [Accessed March 02, 2021].
- [Pajic 2018] Jelena Pajic, José Rivera, Kaiwen Zhang and Hans-Arno Jacobsen. *Eva: Fair and auditable electric vehicle charging service using blockchain*. In Proceedings of the 12th ACM International Conference on Distributed and Event-based Systems, pages 262–265, 2018.
- [Palensky 2011] Peter Palensky and Dietmar Dietrich. *Demand side management: Demand response, intelligent energy systems, and smart loads*. IEEE transactions on industrial informatics, vol. 7, no. 3, pages 381–388, 2011.
- [Parise 2019] Alec Parise, Miguel A. Manso-Callejo, Hung Cao, Marco Mendonca, Harpreet Kohli and Monica Wachowicz. *Indoor Occupancy Prediction using an IoT Platform*. In Proceedings of the 6th International Conference on Internet of Things: Systems, Management and Security, IOTSMS 2019, pages 26–31, Granada, Spain, 2019.
- [Peng 2018] Y. Peng, A. Rysanek, Z. Nagy and A. Schluter. *Using machine learning techniques for occupancy-prediction-based cooling control in office buildings*. Applied Energy, vol. 211, pages 1343–1358, 2018.
- [Pešić 2019] Saša Pešić, Milenko Tošić, Ognjen Iković, Miloš Radovanović, Mirjana Ivanović and Dragan Bošković. *BLEMAT: Data Analytics and Machine Learning for Smart Building Occupancy Detection and Prediction*. International Journal on Artificial Intelligence Tools, vol. 28, 2019.
- [Plappert 2018] Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder et al. *Multi-goal reinforcement learning: Challenging robotics environments and request for research*. arXiv preprint arXiv:1802.09464, 2018.

Bibliography

- [Pourazarm 2014] Sepideh Pourazarm, Christos G. Cassandras and Andreas Malikopoulos. *Optimal routing of electric vehicles in networks with charging nodes: A dynamic programming approach*. In Proceedings of the 2014 IEEE International Electric Vehicle Conference (IEVC), pages 1–7. IEEE, 2014.
- [Priyadarshini 2015] Rashmi Priyadarshini and RM Mehra. *Quantitative review of occupancy detection technologies*. Int. J. Radio Freq, vol. 1, pages 1–19, 2015.
- [Pustišek 2016] Matevž Pustišek, Andrej Kos and Urban Sedlar. *Blockchain based autonomous selection of electric vehicle charging station*. In 2016 international conference on identification, information and knowledge in the Internet of Things (IIKI), pages 217–222. IEEE, 2016.
- [Qolomany 2018] Basheer Qolomany, Ala Al-Fuqaha, Driss Benhaddou and Ajay Gupta. *Role of Deep LSTM Neural Networks and Wi-Fi Networks in Support of Occupancy Prediction in Smart Buildings*. In Proceedings of the IEEE 19th International Conference on High Performance Computing and Communications, HPCC 2017, IEEE 15th International Conference on Smart City, SmartCity 2017 and IEEE 3rd International Conference on Data Science and Systems, DSS 2017, pages 50–57, Bangkok, Thailand, 2018.
- [Rahaman 2019] Mohammad Saiedur Rahaman, Harsh Pare, Jonathan Liono, Flora D. Salim, Yongli Ren, Jeffrey Chan, Shaw Kudo, Tim Rawling and Alex Sinickas. *OccuSpace: Towards a Robust Occupancy Prediction System for Activity Based Workplace*. In Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops 2019, pages 415–418, Kyoto, Japan, 2019.
- [Ramsebner 2021] Jasmine Ramsebner, Reinhard Haas, Amela Ajanovic and Martin Wietschel. *The sector coupling concept: A critical review*. Wiley Interdisciplinary Reviews: Energy and Environment, page e396, 2021.
- [Razavi 2019] Rouzbeh Razavi, Amin Gharipour, Martin Fleury and Ikpe Justice. *Occupancy detection of residential buildings using smart meter data: A large-scale study*. Energy and Buildings, vol. 183, pages 195–208, 2019.
- [Reinsel 2018] David Reinsel, John Gantz and John Rydning. *The Digitization of the World: From Edge to Core*. Data Age 2025, pages 1–28, 2018.
- [Research 2018] Transparency Market Research. *Home automation market - global industry analysis, size, share, growth, trends, and forecast 2018 - 2026*. Market Report, 2018.
- [Ritchie 2019] Hannah Ritchie and Max Roser. *Access to Energy*. Our World in Data, 2019. <https://ourworldindata.org/energy-access>[Accessed January 11, 2021].
- [Roser 2020] Max Roser and Hannah Ritchie. *Technological Progress*. Our World in Data, 2020. <https://ourworldindata.org/technological-progress> [Accessed January 11, 2021].

- [Ryu 2016] Seung Ho Ryu and Hyeun Jun Moon. *Development of an occupancy prediction model using indoor environmental data based on machine learning techniques*. Building and Environment, vol. 107, pages 1–9, 2016.
- [Sadeghianpourhamami 2019] Nasrin Sadeghianpourhamami, Johannes Deleu and Chris Develder. *Definition and evaluation of model-free coordination of electrical vehicle charging with reinforcement learning*. IEEE Transactions on Smart Grid, vol. 11, no. 1, pages 203–214, 2019.
- [Salimi 2019] Shide Salimi and Amin Hammad. *Critical review and research roadmap of office building energy management based on occupancy monitoring*. Energy and Buildings, vol. 182, pages 214–241, 2019.
- [Santamaria 2019] Joseph Santamaria, Belinda Kinkead, Weiping Pan, Peter Stuart and Adam Nancarrow. *Digitalization and the Future of Energy*. , pages 1–28, 2019.
- [Schneider 2014] Michael Schneider, Andreas Stenger and Dominik Goeke. *The electric vehicle-routing problem with time windows and recharging stations*. Transportation Science, vol. 48, no. 4, pages 500–520, 2014.
- [Schoenberg 2019] Sven Schoenberg and Falko Dressler. *Planning ahead for EV: Total travel time optimization for electric vehicles*. In Proceedings of the IEEE Intelligent Transportation Systems Conference, ITSC 2019, pages 3068–3075. IEEE, 2019.
- [Schrittwieser 2020] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepelet *al.* *Mastering atari, go, chess and shogi by planning with a learned model*. Nature, vol. 588, no. 7839, pages 604–609, 2020.
- [Schwab 2016] Klaus Schwab. *The Fourth Industrial Revolution: what it means, how to respond*. World Economic Forum, 2016. <https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond> [Accessed January 11, 2021].
- [Schweiss 2006] Markus Schweiss. *Breakdown of electrical energy consumption in Germany's private households (2006)*. Energy Agency NRW, 2006.
- [Sedlmeir 2020] Johannes Sedlmeir, Hans Ulrich Buhl, Gilbert Fridgen and Robert Keller. *The energy consumption of blockchain technology: beyond myth*. Business & Information Systems Engineering, vol. 62, no. 6, pages 599–608, 2020.
- [SERI 2020] State Secretariat for Education, Research and Innovation SERI. *Digitalization report. Action Plan for Education, Research and Innovation (ERI) 2019–2020*, 2020. <https://www.sbfi.admin.ch/sbfi/en/home/eri-policy/digitalisation.html> [Accessed December 28, 2020].

Bibliography

- [Shin 2020a] Mi Su Shin, Kyu Nam Rhee and Gun Joo Jung. *Optimal heating start and stop control based on the inferred occupancy schedule in a household with radiant floor heating system*. Energy and Buildings, vol. 209, page 109737, 2020.
- [Shin 2020b] Myungjae Shin, Dae Hyun Choi and Joongheon Kim. *Cooperative Management for PV/ESS-Enabled Electric Vehicle Charging Stations: A Multiagent Deep Reinforcement Learning Approach*. IEEE Transactions on Industrial Informatics, vol. 16, pages 3493–3503, 2020.
- [Sociaal en Cultureel Planbureau 2005] Sociaal en Cultureel Planbureau. *Tijdsbestedingsonderzoek 2005 - TBO 2005*. DANS, 2005.
- [Statista 2020] Statista. *Number of electric vehicle charging stations in Switzerland from 2012 to 2020, by type*, 2020. <https://www.statista.com/statistics/933033/number-of-electric-vehicle-charging-stations-in-switzerland/> [Accessed March 24, 2021].
- [Statista 2021] Statista. *Bitcoin average energy consumption per transaction compared to that of VISA as of April 14, 2021*, 2021. <https://www.statista.com/statistics/881541/bitcoin-energy-consumption-transaction-comparison-visa/> [Accessed April 27, 2021].
- [Su 2018] Zhou Su, Yuntao Wang, Qichao Xu, Minrui Fei, Yu-Chu Tian and Ning Zhang. *A secure charging scheme for electric vehicles with smart communities in energy blockchain*. IEEE Internet of Things Journal, vol. 6, no. 3, pages 4601–4613, 2018.
- [Sumic 2020] Zarko Sumic. *Hype Cycle for Digital Grid Transformation Technologies*, 2020. Gartner, Inc., 2020. <https://www.gartner.com/en/documents/3989047/hype-cycle-for-digital-grid-transformation-technologies> [Accessed January 11, 2021].
- [Sweda 2012] T. M. Sweda and D. Klabjan. *Finding minimum-cost paths for electric vehicles*. In Proceedings of the 2012 IEEE International Electric Vehicle Conference, pages 1–4. IEEE, 2012.
- [Szabo 1997] Nick Szabo. *Formalizing and securing relationships on public networks*. First Monday, 1997.
- [Tang 2019] Xiaoying Tang, Suzhi Bi and Ying Jun Angela Zhang. *Distributed routing and charging scheduling optimization for internet of electric vehicles*. IEEE Internet of Things Journal, vol. 6, pages 136–148, 2019.
- [Taylor 2017] S.J. Taylor and B. Letham. *Forecasting at scale*. PeerJ Preprint, vol. 5:e3190v2, 2017.
- [The Ethereum Network & Node Explorer 2021] The Ethereum Network & Node Explorer . *Ethereum Mainnet Statistics*, 2021. <https://www.ethernodes.org> [Accessed February 02, 2021].
- [Uhlenbeck 1930] G. E. Uhlenbeck and L. S. Ornstein. *On the Theory of the Brownian Motion*. Physical Review, vol. 36, pages 823–841, 1930.

- [Umoren 2020] Ifiok Anthony Umoren, Syeda SA Jaffary, Muhammad Zeeshan Shakir, Konstantinos Katzis and Hamed Ahmadi. *Blockchain-Based Energy Trading in Electric Vehicle Enabled Microgrids*. IEEE Consumer Electronics Magazine, 2020.
- [U.S. EIA 2016] Energy Information Agency U.S. EIA. *2012 Commercial Buildings Energy Consumption Survey: Energy Usage Summary*. CBECS 2012, 2016. <https://www.eia.gov/consumption/commercial/reports/2012/energyusage/> [Accessed February 25, 2021].
- [U.S. EIA 2020] Energy Information Administration U.S. EIA. *Energy consumption by sector*. Monthly Energy Review, 2020. https://www.eia.gov/totalenergy/data/monthly/pdf/sec2_3.pdf [Accessed June 3, 2020].
- [Vafeiadis 2017] T. Vafeiadis, S. Zikos, G. Stavropoulos and D. Ioannidis. *Machine learning based occupancy detection via the use of smart meters*. In Proceedings of the International Symposium on Computer Science and Intelligent Controls ISCSIC'17, pages 6–12, Budapest, Hungary, 2017. IEEE.
- [Val d'Herens 2019] Val d'Herens. *Service Green Mobility*, 2019. <https://www.valdherens.ch/en/service-green-mobility-fp45415> [Accessed December 28, 2020].
- [Van Hasselt 2015] Hado Van Hasselt, Arthur Guez and David Silver. *Deep reinforcement learning with double q-learning*. arXiv preprint arXiv:1509.06461, 2015.
- [Vázquez-Canteli 2019] José R Vázquez-Canteli and Zoltán Nagy. *Reinforcement learning for demand response: A review of algorithms and modeling techniques*. Applied Energy, vol. 235, pages 1072–1089, 2019.
- [Vecerik 2017] Mel Vecerik, Todd Hester, Jonathan Scholz, Fumin Wang, Olivier Pietquin, Bilal Piot, Nicolas Heess, Thomas Rothörl, Thomas Lampe and Martin Riedmiller. *Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards*. arXiv preprint arXiv:1707.08817, 2017.
- [Verbeek 2020] Martijn Verbeek. *Why price transparency in eMobility matters*. EVBox, 2020. <https://blog.evbox.com/price-transparency-emobility> [Accessed March 02, 2021].
- [Voss 2021] Alex Voss. *How to bring GDPR into the digital age*. Politico, 2021. <https://www.politico.eu/article/gdpr-reform-digital-innovation/> [Accessed April 16, 2021].
- [Waldrop 2016] M. Mitchell Waldrop. *The chips are down for Moore's law*. Nature News, vol. 530, no. 7589, pages 1–144, 2016.
- [Wan 2018a] Zhiqiang Wan, Hepeng Li Haibo He and Danil Prokhorov. *Model-Free Real-Time EV Charging Scheduling Based on Deep Reinforcement Learning*. IEEE Transactions on Smart Grid, vol. 10, pages 5246–5257, 2018.

Bibliography

- [Wan 2018b] Zhiqiang Wan, Hepeng Li and Haibo He. *Residential energy management with deep reinforcement learning*. In 2018 International Joint Conference on Neural Networks (IJCNN), pages 1–7, Rio de Janeiro, Brazil, 2018. IEEE.
- [Wang 2008] Shengwei Wang and Zhenjun Ma. *Supervisory and optimal control of building HVAC systems: A review*. *Hvac&R Research*, vol. 14, no. 1, pages 3–32, 2008.
- [Wang 2018] Yi Wang, Qixin Chen, Tao Hong and Chongqing Kang. *Review of Smart Meter Data Analytics: Applications, Methodologies, and Challenges*. *IEEE Transactions on Smart Grid*, vol. 10, pages 3125–3148, 2018.
- [Wang 2019a] Wei Wang, Tianzhen Hong, Ning Xu, Xiaodong Xu, Jiayu Chen and Xiaofang Shan. *Cross-source sensing data fusion for building occupancy prediction with adaptive lasso feature filtering*. *Building and Environment*, vol. 162, page 106280, 2019.
- [Wang 2019b] Yuntao Wang, Zhou Su and Ning Zhang. *BSIS: Blockchain-based secure incentive scheme for energy delivery in vehicular energy network*. *IEEE Transactions on Industrial Informatics*, vol. 15, no. 6, pages 3620–3631, 2019.
- [Wang 2019c] Zhe Wang, Tianzhen Hong, Mary Ann Piette and Marco Pritoni. *Inferring occupant counts from Wi-Fi data in buildings through machine learning*. *Building and Environment*, vol. 158, pages 281–294, 2019.
- [Weidman 2018] Louia Weidman. *Side looking occupancy sensor*, January 2018. US Patent 9,867,259.
- [WHO 2021] World Health Organization WHO. *Coronavirus disease (COVID-19)*, 2021. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019> [Accessed February 25, 2021].
- [Wilhelm 2021] Sebastian Wilhelm, Dietmar Jakob, Jakob Kasbauer and Diane Ahrens. *GeLaP: German Labeled Dataset for Power Consumption*. In , 02 2021.
- [Xiong 2018] Jiechao Xiong, Qing Wang, Zhuoran Yang, Peng Sun, Lei Han, Yang Zheng, Haobo Fu, Tong Zhang, Ji Liu and Han Liu. *Parametrized deep q-networks learning: Reinforcement learning with discrete-continuous hybrid action space*. arXiv preprint arXiv:1810.06394, 2018.
- [Yan 2015] Da Yan, William O’Brien, Tianzhen Hong, Xiaohang Feng, H Burak Gunay, Farhang Tahmasebi and Ardeshtir Mahdavi. *Occupant behavior modeling for building performance simulation: Current state and future challenges*. *Energy and Buildings*, vol. 107, pages 264–278, 2015.
- [Yang 2012] Zheng Yang, Nan Li, Burcin Becerik-Gerber and Michael Orosz. *A multi-sensor based occupancy estimation model for supporting demand driven HVAC operations*. In *Proceedings of the 2012 Symposium on Simulation for Architecture and Urban Design*, pages 1–8, Orlando, USA, 2012.

- [YCharts 2021] YCharts. *Ethereum Average Transaction Fee as of April 26, 2021*, 2021. https://ycharts.com/indicators/ethereum_average_transaction_fee [Accessed April 27, 2021].
- [Ye 2020] Yujian Ye, Dawei Qiu, Xiaodong Wu, Goran Strbac and Jonathan Ward. *Model-Free Real-Time Autonomous Control for A Residential Multi-Energy System Using Deep Reinforcement Learning*. IEEE Transactions on Smart Grid, vol. 11, no. 4, pages 3068–3082, 2020.
- [Yu 2019] Liang Yu, Weiwei Xie, Di Xie, Yulong Zou, Dengyin Zhang, Zhixin Sun, Linghua Zhang, Yue Zhang and Tao Jiang. *Deep Reinforcement Learning for Smart Home Energy Management*. IEEE Internet of Things Journal, vol. 7, no. 4, pages 2751–2762, 2019.
- [Yu 2020] Liang Yu, Yi Sun, Zhanbo Xu, Chao Shen, Dong Yue, Tao Jiang and Xiaohong Guan. *Multi-agent deep reinforcement learning for HVAC control in commercial buildings*. IEEE Transactions on Smart Grid, vol. 12, no. 1, pages 407–419, 2020.
- [Zhang 2018] Tianyang Zhang, Himanshu Pota, Chi-Cheng Chu and Rajit Gadh. *Real-time renewable energy incentive system for electric vehicles using prioritization and cryptocurrency*. Applied energy, vol. 226, pages 582–594, 2018.
- [Zhang 2020a] Cong Zhang, Yuanan Liu, Fan Wu, Bihua Tang and Wenhao Fan. *Effective charging planning based on deep reinforcement learning for electric vehicles*. IEEE Transactions on Intelligent Transportation Systems, vol. June, pages 1–13, 2020.
- [Zhang 2020b] Feiye Zhang, Qingyu Yang and Dou An. *CDDPG: A Deep Reinforcement Learning-Based Approach for Electric Vehicle Charging Control*. IEEE Internet of Things Journal, 2020.
- [Zhang 2020c] Hongming Zhang and Tianyang Yu. *Taxonomy of Reinforcement Learning Algorithms*. In Deep Reinforcement Learning, pages 125–133. Springer, 2020.
- [Zhang 2020d] Qian Zhang, Kui Wu and Yang Shi. *Route planning and power management for PHEVs with reinforcement learning*. IEEE Transactions on Vehicular Technology, vol. 69, pages 4751–4762, 2020.
- [Zhao 2015] Yang Zhao, Wim Zeiler, Gert Boxem and Timi Labeodan. *Virtual occupancy sensors for real-time occupancy information in buildings*. Building and Environment, vol. 93, pages 9–20, 2015.
- [Zhao 2018] Bochao Zhao, Lina Stankovic and Vladimir Stankovic. *Electricity usage profile disaggregation of hourly smart meter data*. 4th International Workshop on Non-Intrusive Load Monitoring, March 2018.
- [Zhou 1997] K. Zhou and S. Yang. *The Neurothermostat: predictive optimal control of residential heating systems*. Advances in Neural Information Processing Systems, vol. 9, pages 953–959, 1997.

Bibliography

- [Zhou 2016] Kailie Zhou and Shanlin Yang. *Understanding household energy consumption behavior: the contribution of energy big data analytics*. Renewable and Sustainable Energy Reviews, vol. 56, pages 810–819, 2016.
- [Zhou 2019] Zhenyu Zhou, Bingchen Wang, Yufei Guo and Yan Zhang. *Blockchain and computational intelligence inspired incentive-compatible demand response in internet of electric vehicles*. IEEE Transactions on Emerging Topics in Computational Intelligence, vol. 3, no. 3, pages 205–216, 2019.
- [Zimmerman 2018] L. Zimmerman, R. Weigel and G. Fischer. *Fusion of non-intrusive environmental sensors for occupancy detection in smart homes*. IEEE Internet of Things, vol. 5, pages 2343–2352, 2018.
- [Zoha 2013] A. Zoha, A. Gluhak, M. Nati and M. A. Imran. *Low-power appliance monitoring using Factorial Hidden Markov Models*. In 2013 IEEE Eighth International Conference on Intelligent Sensors, Sensor Networks and Information Processing, pages 527–532. IEEE, April 2013.
- [Zuo 2017] Sixiang Zuo, Zhiyang Wang, Xiaorui Zhu and Yongsheng Ou. *Continuous reinforcement learning from human demonstrations with integrated experience replay for autonomous driving*. In 2017 IEEE International Conference on Robotics and Biomimetics (ROBIO), pages 2450–2455, Macau SAR, China, 2017. IEEE.
- [Zwanzger 2020] Freddy Zwanzger. *Energy Web Chain – Smart contract cluster analytics*, 2020. <https://www.anyblockanalytics.com/blog/energy-web-chain-smart-contract-cluster-analytics/> [Accessed February 02, 2021].

Marina Dorokhova

Avenue de la Confrérie 44, 17, 1008 Prilly, Switzerland | 21.10.1993

☎ +41 78 674 88 12 | ✉ marina.dorokhova@epfl.ch | 🌐 marina-dorokhova

Education

EPFL (Ecole Polytechnique Fédérale de Lausanne)

Neuchâtel, Switzerland

PhD Candidate in Doctoral School of Energy, Laboratory of Photovoltaics

Nov. 2017 - Jul. 2021

- Team leader in energy monitoring, forecasting and optimization in EU Horizon 2020 project "FEEdBACK"
- Researcher in SCCER-FURIES project "Blockchain exchanges for green mobility"

Global Young Scientists Summit - Singapore	top 5%	2020
Next Generation Women Leaders by McKinsey	top 1.5%	2019
Business Concept start-up training by InnoSuisse	top 5%	2018

EPFL (Ecole Polytechnique Fédérale de Lausanne) 5.7/6.0

Lausanne, Switzerland

MSc. in Energy Management & Sustainability

Sept. 2015 - Oct. 2017

- Event Manager and President at Energy Business Policy and Technology Group (EBPTG) at EPFL
- "Design of energy management system for robots operating in harsh and hazardous environment" *Master thesis, CERN*

Best GPA prize	top 5%	2017
MES Fellowship	top 5%	2015

BMSTU (Bauman Moscow State Technical University) 5.0/5.0

Moscow, Russia

BSc. in Power Engineering

Sept. 2011 - Jun. 2015

- Gas turbines and non-conventional power plants department
- "Design of wind-powered energy refrigeration unit for temporary storage of agricultural crops" *Bachelor thesis, BMSTU*

Silver Sign Award for diploma with honors	top 1%	2015
FONDRE scholarship for young entrepreneurs	top 0.5%	2014
Academic Council scientific scholarship	top 1%	2014
Russian President's Scholar	top 1%	2013
V. Potanin Foundation scholarship	top 3%	2013
"Lift to the future" scholarship	top 5%	2013
Siemens Science Award	III place	2011

Experience

CERN (European Organization for Nuclear Research)

Geneva, Switzerland

Engineering Intern in EN-STI-ECE Robotics

Feb. 2017 - Aug. 2017

- Designed robot's energy management system based on complete power and locomotion models. Proposed solution decreased the battery size by half and achieved 7% energy savings during robotic missions. The algorithms were developed in C++.
- Produced software has been integrated into robot's functional framework thus contributing to its secure operation in harsh environment of CERN accelerators.

First Solar Inc.

Dubai, UAE

Business Development Analyst Intern

Jul. 2016 - Sept. 2016

- Conducted performance review and evaluation of the Phase I of the Mohamed bin Rashid Al Maktoum Solar Park. Improved the accuracy of energy predictions on more than 5 different aspects. The analysis was introduced to the top-management and included in company's guidelines for simulations.
- Realized a joint project with General Electric on modelling and economical assessment of the PV-RO desalination plants in UAE region. Proposed best-case scenario and 4 different system configurations.

Energy Business, Policy and Technology Group

Lausanne, Switzerland

Event Manager and President

Sept. 2015 - Apr. 2017

- Organized scientific conference at EPFL with 50 participants, involved 5 speakers from top Swiss companies and universities.
- Established annual student energy conference, 50+ participants from 9 different countries. Attracted speakers from Forbes 30 under 30 list and 8 international companies. Co-Founded international organization for energy students and young professionals in Europe "YES-Europe".

Additional Information

- **Research interests:** Energy Management | Intelligent Energy Systems | Energy Informatics | Integration of Renewables | Building Automation | Machine Learning | Optimization | Forecasting | Data Science | Blockchain
- **IT skills:** Microsoft Office | Mac OS | MATLAB | Python | C++ | AMPL | LaTeX | Autocad | Siemens NX | ANSYS | LabVIEW | PVsyst
- **Languages:** Russian (Native) | English (C2) | French (C2) | Italian (B2)
- **Hobbies:** Master of sports in ballroom dancing

Publications

- **Real-world implementation of an ICT-based platform to promote energy efficiency**
Marina Dorokhova, Fernando Ribeiro, António Barbosa, João Viana, Filipe Soares, Nicolas Wyrsh
[Energies, 14\(9\), 2416, 2021](#)
- **FEEDBACK: An ICT-based platform to increase energy efficiency through buildings' consumer engagement**
Filipe Soares, André Madureira, Andreu Pagès, António Barbosa, António Coelho, Fernando Cassola, Fernando Ribeiro, João Viana, José Ricardo Andrade, Marina Dorokhova, Nelson Morais, Nicolas Wyrsh, Trine Sørensen
[Energies, 14\(6\), 1524, 2021](#)
- **Deep reinforcement learning control of electric vehicle charging in the presence of photovoltaic generation**
Marina Dorokhova, Yann Martinson, Christophe Ballif and Nicolas Wyrsh
[Submitted, 2021](#)
- **A blockchain-supported framework for charging management of electric vehicles**
Marina Dorokhova, Jérémie Vianin, Jean-Marie Alder, Christophe Ballif, Nicolas Wyrsh and David Wannier
[Submitted, 2021](#)
- **Routing of electric vehicles with intermediary charging stations: A reinforcement learning approach**
Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh
[Frontiers in Big Data, 4:586481, 2021](#)
- **Behavioral Change towards EE by Utilizing ICT Tools**
Albert Hoffrichter, Evangelos Zacharis, Angelina Katsifaraki, Ashley Morton, Gloria Calleja, Federica Fuligni, Marko Batic, Marina Dorokhova, Niall Castelli, Konstantinos Kanellos and Thanh Nguyen
[MDPI Proceedings of The 8th Annual International Sustainable Places Conference \(SP2020\), vol. 65, no. 1: 10, 2020](#)
- **Rule-based scheduling of air conditioning using occupancy forecasting**
Marina Dorokhova, Christophe Ballif and Nicolas Wyrsh
[Energy and AI, vol. 2, page 100022, 2020](#)
- **Unsupervised algorithm for disaggregating low-sampling-rate electricity consumption of households**
Jordan Holweger, Marina Dorokhova, Lionel Bloch, Christophe Ballif and Nicolas Wyrsh
[Sustainable Energy, Grids and Networks, vol. 19, page 100244, 2019](#)
- **Energy refrigeration unit based on the Andreau wind turbine**
Marina Dorokhova
[Youth Scientific-Technical Journal vol.4, 2015](#)
- **Organic Rankine Cycle and its application in renewable energy**
Belov Georgiy, Marina Dorokhova
[Engineering Bulletin vol.2, 2014](#)
- **Chute's calculation for feeding parts to industrial robots acquisition position**
Marina Dorokhova
[Youth Scientific-Technical Journal vol.3, 2013](#)

Certifications

- **Improving Deep Neural Networks: Hyperparameter tuning, Regularization and Optimization**
Credential ID: LL4XY4JEL4BC
[Coursera, March 2020](#)
- **Neural Networks and Deep Learning**
Credential ID: 4T64RT582VLS
[Coursera, March 2020](#)
- **Blockchain Specialization**
Credential ID: A52JFAVJGGAM
[Coursera, April 2019](#)