INSTITUTO SUPERIOR DE ENGENHARIA DO PORTO

MESTRADO EM ENGENHARIA DE INTELIGÊNCIA ARTIFICIAL




Gestão Inteligente e Distribuída de Comunidades de Cidadãos

RAFAEL DUARTE PEREIRA DA SILVA

Maio de 2025

Distributed Intelligent Management of Citizen Communities

Rafael Duarte Pereira da Silva

1190979

Dissertation to obtain the Master's Degree in

Artificial Intelligence Engineering

Supervisor: Doctor Luis Gomes, Assistant Professor, Polytechnic of Porto - School of Engineering

Co-supervisor: Doctor Zita Vale, Associate Professor, Polytechnic of Porto - School of Engineering

Júri:

President:

Carlos Fernando da Silva Ramos; Full Professor, Polytechnic of Porto - School of Engineering

Vocals:

Luiz Felipe Rocha de Faria; Associate Professor, Polytechnic of Porto - School of Engineering

Luís Filipe de Oliveira Gomes; Assistant Professor, Polytechnic of Porto - School of Engineering

Porto, May 2025

DECLARAÇÃO DE INTEGRIDADE

DECLARAÇÃO DE INTEGRIDADE

Declaro ter conduzido este trabalho académico com integridade. Não plagiei ou apliquei qualquer forma de uso indevido de informações ou falsificação de resultados ao longo do processo que levou à sua elaboração.

Declaro que o trabalho apresentado neste documento é original e de minha autoria, não tendo sido utilizado anteriormente para nenhum outro fim.

Declaro ainda que tenho pleno conhecimento do Código de Conduta Ética do P.PORTO.

ISEP, Porto, 19 de maio de 2025

Rafael Silva

Resumo

A utilização e integração de modelos inteligentes nos edifícios pode transformar as experiências dos utilizadores dentro do edifício, proporcionando a otimização dos espaços e formas eficientes de utilizar e interagir com os recursos do edifício. A utilização de soluções inteligentes traz alguns desafios que devem ser estudados, como a heterogeneidade entre os recursos e a necessidade de adaptar os edifícios já existentes ao conceito de edifícios inteligentes.

Embora os edifícios inteligentes possam revolucionar a forma como as pessoas utilizam e interagem com os espaços, o grupo de edifícios, ao criar comunidades, traz novas oportunidades para permitir que os membros interligados atinjam objetivos comuns, modelando papéis cooperativos, colaborativos e, por vezes, competitivos. Esta nova dinâmica em que os sistemas orgânicos podem comunicar e interagir também levanta desafios quanto à modelação dos utilizadores, às suas preferências e à existência de infraestruturas comuns para permitir a implementação de modelos inteligentes ao nível da comunidade, edifício e utilizador.

Esta dissertação tem como objetivo conceber, implementar, testar e validar uma infraestrutura baseada em *containers*, intitulada Caravels, que combina os conceitos de comunidades inteligentes e edifícios inteligentes para desenvolver uma solução sensível ao contexto que considera diferentes utilizadores e edifícios. A solução concebida emprega uma arquitetura distribuída para a gestão de comunidades inteligentes de cidadãos, onde cada membro opera como uma entidade autónoma, enquanto permanece interligado através de uma infraestrutura partilhada. A arquitetura permite serviços tanto a nível local como comunitário, sendo que um membro pode integrar serviços individuais, escolhidos especificamente para esse utilizador, ao mesmo tempo que contribui e beneficia de otimizações a nível comunitário.

Central ao projeto está a modelação das preferências do utilizador em ambientes complexos, dinâmicos e multiutilizador. A dissertação explora os desafios psicológicos e cognitivos da representação de preferências, reconhecendo que os utilizadores têm dificuldades em articular ou priorizar as suas próprias preferências. Os modelos propostos podem adaptar-se ao longo do tempo, incorporando *feedback* e dados comportamentais para apoiar a tomada de decisões proativas e conscientes do contexto. As técnicas de inteligência artificial, incluindo a aprendizagem supervisionada, não supervisionada e por reforço, estão integradas em todo o sistema para permitir a análise preditiva, a otimização e o controlo autónomo.

Para validar a arquitetura e as metodologias propostas, foram conduzidos vários estudos de caso em cenários realistas, refletindo as diferentes necessidades dos utilizadores, procura de energia e recursos distribuídos. Os resultados demonstram que o sistema pode modelar o comportamento do utilizador, apoiar a cooperação a nível comunitário e melhorar a eficiência e a inteligência geral do edifício inteligente. Os resultados desta dissertação contribuíram para seis publicações científicas, incluindo uma revista com um fator de impacto de 6,6.

Palavras-chave: arquitetura de *containers*, computação distribuída, comunidades inteligentes, edifícios inteligentes, modelos inteligentes, modelação de preferências do utilizador

Abstract

The use and integration of intelligent models in buildings can transform the users' experiences inside the building, being able to provide optimization of spaces and efficient ways to use and interact with the building resources. The use of these intelligent solutions brings some challenges that must be studied, such as the heterogeneity among resources and the need of retrofitting to update the already existing buildings to the concept of intelligent buildings.

Although intelligent buildings can revolutionize how people use and interact with spaces, the collective of buildings, creating communities, brings new opportunities to enable interconnected members to pursue common goals, modeling cooperative, collaborative, and, sometimes, competitive roles. This new dynamic where organic systems can communicate and interact also raises challenges regarding the modeling of users, users' preferences, and the existence of common infrastructures to enable the implementation of intelligent models at community, building, and user levels.

This dissertation aims to conceive, implement, test, and validate a container-based infrastructure, entitled Caravels, which combines the concepts of intelligent communities and intelligent buildings to develop a context-aware solution that considers different users and different buildings. The conceived solution, Caravels, employs a distributed architecture for the intelligent management of citizen intelligent communities, where each member operates as an autonomous and intelligent entity, while remaining interconnected through a shared infrastructure. The architecture enables both local and community-level services, meaning that a member can integrate individual services, specifically chosen for that user, while still contributing to and benefiting from community-level optimizations.

Central to this work is the modeling of user preferences in complex, dynamic, and multi-user environments. The dissertation explores the psychological and cognitive challenges of preference representation, recognizing that users often struggle to articulate or prioritize their own preferences. The proposed models can adapt over time, incorporating feedback and behavioral data to support proactive and context-aware decision-making. Artificial intelligence techniques, including supervised, unsupervised, and reinforcement learning, are integrated across the system to enable predictive analytics, optimization, and autonomous control.

To validate the proposed architecture and methodologies, several case studies were conducted in realistic scenarios, reflecting heterogeneous user needs, fluctuating energy demands, and distributed resources. The results demonstrate that the system can effectively model user behavior, support community-level cooperation, and enhance the overall efficiency and intelligence of the intelligent building. The results of this dissertation contributed to six scientific publications, including one journal with an impact factor of 6.6.

Keywords: container-based architecture, distributed computing, intelligent communities, intelligent buildings, intelligent models, user modeling, user preferences.

Index

List of Figure

List of Tables

List of Acronyms

AI	Artificial Intelligence
API	Application Programing Interface
BLE	Bluetooth Low Energy
CI/CD	Continuous Integration/Continuous Deployment
CNN	Convolutional Neural Network
CPU	Central Processing Unit
DR	Demand Response
EMS	Energy Management System
ESS	Energy Storage System
EU	European Union
EUR	Euro
EV	Electric Vehicle
FRP	Fast Reverse Proxy
GDPR	General Data Protection Regulation
GECAD	Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development
GNN	Graph Neural Network
GPU	Graphics Processing Unit
H2H	Human-to-Human
HAR	Human Action-Reaction
HTTP	Hypertext Transfer Protocol
HVAC	Heating, Ventilation and Air Conditioning
IF	Impact Factor
IoT	Internet of Things
MDP	Markov Decision Process
ML	Machine Learning
MQTT	Message Queuing Telemetry Transport
NGS	New Generation Storage
OS	Operating System
P2P	Peer to Peer
PV	Photovoltaic

R&D	Research and Development
REST	Representational State Transfer
RL	Reinforcement Learning
RNN	Recurrent Neural Network
RQ	Research Question
UC	Use Case
UI	User Interface
VM	Virtual Machine
VPN	Virtual Private Network

1 Introduction

This initial chapter provides a general contextualization of the work described in this dissertation. This chapter presents the contextualization of the work, the main research questions and objectives, the scientific contributions, and the document organization.

1.1 Contextualization

The increase in greenhouse gas emissions, especially due to the higher energy consumption, and the increase in population led to the search of greener, efficient, and more intelligent solutions for the communities. Knowing that, buildings are responsible for 40% of energy consumption and 36% of greenhouse gas emissions (*Energy Saving: EU Action to Reduce Energy Consumption | News | European Parliament*, n.d.), one of the objectives proposed by the European Parliament is that, by 2050, half of Europe's citizens could be producing up to half of the European Union's (EU) renewable energy (*Energy Communities*, n.d.).

In (DIRECTIVE (EU) 2018/2001, 2018), the EU has defined energy communities as "legal entities that are constituted by citizens, local authorities, businesses or other legal entities with a common interest in producing, consuming or saving energy collectively in a specific geographical area." Energy communities are seen as a key element of the EU's strategy to decarbonize its economy and promote energy efficiency. They can provide a range of benefits such as distributed generation of renewable energy, improved energy efficiency, local control of energy resources (*Energy Communities Repository - General Information - European Commission*, n.d.). The number of energy communities in Europe is growing rapidly. In 2023, there were over 9,000 energy communities, in total, working towards reducing the EU's annual €330 billion bill for energy imports (*Energy Saving: EU Action to Reduce Energy Consumption | News | European Parliament*, n.d.).

While energy communities gathered significant attention for their potential in optimizing energy utilization and promoting sustainability, a majority of research has focused solely on the energy sector, overlooking the vast potential for resource sharing to extend to other domains, including transportation, food, water management or even data. By expanding the scope of resource sharing, communities can achieve greater efficiency, sustainability, and resilience across a wide range of resource-intensive activities.

In the context of communities, it is often necessary to concede individual objectives for the greater good of the community, as individual actions can significantly impact the overall well-being and sustainability of the community. Communities that embrace distributed solutions, involving multiple stakeholders, pursuing multiple objectives are well-positioned to address the challenges and seize the opportunities (Jakobsen et al., 2023). Although only considering energy domain, there are solutions that compromise in dealing with the importance of considering individual versus collective (Coignard et al., 2023) and others with optimizing for multiple objectives (Xiong et al., 2019; Z. Yang et al., 2021).

Giving that communities are composed by buildings of different construction years, older ones should not be discarded knowing that it is estimated 85–95% of the buildings that exist today will be standing in 2050 (Maduta et al., 2022). For older ones, smart retrofit can significantly impact energy efficiency and carbon emissions (Hashempour et al., 2019). Building energy retrofitting can include a variety of energy efficiency measures, such as improving the envelope by incorporating insulation, energy-efficient windows, airtightness, the use of energy-efficient lighting systems and high-performance heating, ventilation and air conditioning (HVAC) systems (Ibaseta et al., 2021a; Peiris et al., 2023).

Protecting individual privacy in intelligent communities is paramount to fostering trust and ensuring the well-being of community members. The mitigation of privacy risks can be achieved by adopting a rights-based approach ensuring that individuals have control over their personal data (Y. T. Lee et al., 2017) and minimizing the collection and processing of sensitive information (Lin et al., 2022). Intelligent communities are also vulnerable to cyberattacks, as they rely on interconnected networks and store valuable data. Robust cybersecurity measures are essential to protect against intrusions, data breaches, and other malicious activities. This increases the need for comprehensive cybersecurity strategies in smart communities, encompassing data encryption, authentication protocols, and intrusion detection systems (Ismagilova et al., 2020). A balanced approach that integrates privacy-enhancing technologies, rights-based frameworks, and robust cybersecurity measures is crucial for creating trustworthy and resilient smart communities.

However, existing solutions often overlook the occurrence of anomalous or unexpected situations. This highlights the need for further research into solutions that consider the broader environmental context and the organization of effective mitigation strategies. Additionally, there is a pressing need to investigate synergistic improvements in multi-purpose energy communities, exploring ways to optimize resource utilization and maximize benefits across various domains (de São José et al., 2021).

A robust and adaptive user preference modeling framework plays a pivotal role in advancing intelligent community management solutions (To et al., 2018). By accurately capturing and interpreting diverse user needs, habits, and contextual cues, such a framework enables systems to personalize services and optimize interactions at both individual and community levels. Moreover, preference models must be capable of adapting to evolving user behaviors and heterogeneous environments, ensuring that personalization remains relevant and effective over time (Dasgupta et al., 2019; Johnston et al., 2019). Effective preference modeling empowers community members to engage more actively with intelligent systems, fostering a sense of inclusion and responsiveness. Crucially, these models must be designed with accessibility and inclusivity in mind, ensuring that they reflect the needs of all users—regardless of their technical literacy, cognitive abilities, or engagement levels.

1.2 Problem Statement

This thesis addresses the problem of the effective management of citizen communities which is crucial for ensuring their sustainability, efficiency, and overall well-being. Traditional approaches to community management often lack the sophistication and integration needed to address the complex challenges faced by modern communities.

Current community management practices typically focus on individual domains, such as energy optimization, waste management, or mobility. This siloed approach fails to capture the interconnectedness of these domains and overlooks opportunities for optimizing resource utilization and improving community outcomes.

Moreover, existing community management solutions often lack flexibility and adaptability to cater to the diverse needs and preferences of community members. Traditionally, the user interfaces of these solutions often require users to navigate complex menus and systems, making it difficult for them to effectively engage with and manage their communities.

In addition, data security and privacy concerns pose significant challenges to the adoption of intelligent community management solutions. The collection, storage, and analysis of community data raise concerns about unauthorized access and misuse of personal information.

1.3 Objectives and Research Questions

The objective of this thesis is to conceive, develop, test and validate a holistic and practical solution for multi-domain distributed intelligent community management, leveraging advanced technologies to enhance collaboration, resource optimization, and sustainability. This main objective has been decomposed in the following objectives:

- O1 - To design and develop a multi-domain solution for context-aware intelligent community management supported by IoT devices in retrofitted buildings;
- O2 - To investigate and implement artificial intelligence techniques, such as machine learning and optimization algorithms, for optimizing community resources considering the users preferences and needs;

- O3 - To explore the development of dynamic and adaptive user modeling systems to enhance user interaction and provide personalized support to community members;
- O4 - To develop a container-based distributed architecture to support the seamless integration of multiple management models in intelligent communities;
- O5 - To test and evaluate the feasibility, effectiveness, efficiency and impact of the proposed solution through real-world case studies and user testing with a diverse range of community members.

To better guide the research within the scope of this work, four primary research questions were formulated and are described as follows:

- *RQ1 - Can a multi-domain solution for intelligent community management be developed and implemented using IoT devices, artificial intelligence, secure information sharing, and distributed computing?*
- *RQ2 - Can artificial intelligence models be effectively utilized to optimize resource usage, share useful data, consider user needs, improve service mobility, and enhance social engagement within citizen communities?*
- *RQ3 - Can dynamic and adaptive user modeling be designed to provide personalized and user-friendly interactions for managing citizen communities?*
- *RQ4 - What technological tools could be applied to create a secure and distributed infrastructure for containers in heterogeneous machines?*

1.4 Scientific Contributions

This work was developed in the Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development (GECAD)[1], which has as its mission the promotion and development of scientific research in the knowledge and decision sciences domains, with Information Technologies as support. It involves two main areas: intelligent systems and power energy systems.

The work developed in this dissertation was supported by six R&D projects funded by FCT, P2020, P2030, PRR:

- TioCPS[2] - Trustworthy and Smart Communities of Cyber-Physical Systems (POCI-01-0247-FEDER-046182);
- ITEA TioCPS[3] – Trustworthy and Smart Communities of Cyber-Physical Systems (18008 TIoCPS);
- NGS[4] - New Generation Storage (PRR 02-C05-i01.01-2022.PC644936001-00000045);

[1] GECAD - https://www.gecad.isep.ipp.pt/
[2] TIoCPS - https://www.tiocps.pt/
[3] ITEA TioCPS - https://itea4.org/project/tiocps.html/
[4] NGS - https://newgenstorage.pt/

- EnergyP2P [5] – *Comercialização de Energia baseado em* Peer-to-Peer trading (COMPETE2030-FEDER-00960500–14646);
- Sa4CPS [6] - Secure situational awareness for critical cyber-physical systems (COMPETE2030-FEDER-00368800-12737);
- ITEA Sa4CPS[7] - Secure situational awareness for critical cyber-physical systems (22007 Sa4CPS).

The TIoCPS project aims to contribute towards solving this grand challenge, boosting the businesses of the involved industries, especially in the context of selected use cases (UCs) related to energy and mobility-related industrial application sectors for enabling a more smart, interoperable, and sustainable ecosystem and society by enabling trustworthy and smart communities for CPS. For this project, Caravels was used as a framework to deploy all the developed services, including user interfaces for each building and for the community operator. Also, during this project were containerized forecast models for energy consumption, generation, and flexibility. Additionally, the integration and containerization of a demand response program were implemented, which involved communication with each building and its IoT devices.

The NGS project aim is to structurally transform the national production fabric, creating the necessary conditions, in terms of technology and human resources, for an industrial ecosystem capable of mass production of innovative technologies, and a complete value chain that allows for world-class end-of-life management. In this project, Caravels was utilized to deploy energy management services, which collected data from forecasting services. Additionally, scheduled operations were integrated into a user interface to facilitate the interpretation of results.

Caravels will serve as a development and deployment framework for pilot projects. In the EnergyP2P project, it will focus on energy storage management models and peer-to-peer (P2P) energy transactions. Meanwhile, in the Sa4CPS project, Caravels will implement critical operational mechanisms, such as a distribution network for an energy community.

Throughout these projects several prototypes were developed resulting in a total of four prototypes:

- Energy Community Simulation Software;
- Building Management System;
- Community Management System;
- Distributed Management System for Energy Communities.

Besides the direct contribution to these scientific projects, throughout the development of this work, a total of four scientific papers were published, from which one was published in a scientific journal and the other in scientific conferences and book chapters:

[5] EnergyP2P - https://www.gecad.isep.ipp.pt/portfolio/_energyp2p/
[6] Sa4CPS - https://sa4cps.pt/
[7] ITEA Sa4CPS - https://itea4.org/project/sa4cps.html

- [Book Chapter] **Rafael Silva**, Bruno Ribeiro, Luis Gomes, Zita Vale "Home Energy Management Models", accepted in "Home Digital Twins";
- [Conference] **Rafael Silva**, Rita Costa, Luis Gomes, Zita Vale "CARAVELS: a case study using peer-to-peer transaction, voluntary demand response, and energy storage optimization", accepted in "21st European Energy Market Conference" (EEM25);
- [Journal – IF 6.6] Rita Costa, **Rafael Silva**, Ricardo Faia, Luis Gomes, Pedro Faria, Zita Vale (2024) "Empowering Energy Communities: A Comprehensive Study on Distributed Energy Storage Systems for Sustainable Consumption", published in Energy & Buildings in 2024 doi: 10.1016/j.enbuild.2024.114953 – in this paper the author of this thesis contributed to the development, containerization, and deployment of all data services and models;
- [Conference] Bruno Ribeiro, **Rafael Silva**, Bruno Mota, Luis Gomes, Zita Vale "Learning-Based Models for Intelligent Control Over Air Conditioning Units in a Smart Building", presented in "19th International Conference on Soft Computing Models in Industrial and Environmental Applications" (SOCO 2024), doi: 10.1007/978-3-031-75013-7_19 – the author of this dissertation contributed to the development of one of the models and the integration within a real system;
- [Conference] **Rafael Silva**, Luis Gomes, Zita Vale "Caravels: a Decentralized Container-Based Infrastructure for Sustainable Human-Centric Intelligent Energy Communities", presented in "19th International Conference on Soft Computing Models in Industrial and Environmental Applications" (SOCO 2024), doi: 10.1007/978-3-031-75013-7_23 – this paper received an invitation to be submitted to Neurocomputing journal;
- [Conference] Bruno Ribeiro, **Rafael Silva**, Luis Gomes, Zita Vale "Detailed analysis of a deep learning energy forecast model considering different input units and magnitudes", presented in 12th IFAC Symposium on Control of Power and Energy Systems - CPES 2024, doi: 10.1016/j.ifacol.2024.07.500 - the author of this dissertation contributed to the integration of the model within a platform, which enabled visualization of the results.

In addition to the scientific papers already accepted and published, another paper is submitted to journals, currently under review:

- [Journal – IF 5.5 – Under review] **Rafael Silva**, Luis Gomes, Zita Vale "Distributed Computing for Intelligent Buildings: The Caravels Approach", Neurocomputing.

In addition to these publications, the author of this dissertation presented part of the papers at the international conference:

- 19th International Conference on Soft Computing Models in Industrial and Environmental Applications (SOCO 2024) – Salamanca, Spain.

From this master thesis three bachelor projects, from Electrical and Computer Engineering, were supervised:

- “Single-board computers configuration using linux-based immutable operating systems”, finished with 16 of 20;
- “Open Source IoT devices for Intelligent Buildings”, finished with 17 of 20;
- “IoT device for centralization of notifications”, still ongoing.

1.5 Document Structure

This dissertation is divided into six chapters, each has been prepared and organized to facilitate reading, being read in its entirety or selected chapters.

The first chapter introduces and describes the contextualization for the development of the work presented in the dissertation. Additionally, the objectives and research questions of this dissertation are outlined, as well as its scientific contributions.

In the second chapter, the state of the art of the domains covered by this work is discussed, namely: intelligent communities, intelligent buildings, user preference modeling, and finally, artificial intelligence.

The third chapter describes the methods, models, materials, and tools used in the development of the work presented in this dissertation. Besides that, it contextualizes this work towards containerization and orchestration, distributed systems, data privacy and security, and possible ethical and social issues associated with it.

The framework developed in the scope of this work, entitled Caravels, is presented in the fourth chapter. Firstly, the architecture is explained and then the way on how containers are integrated is presented. Secondly, the services and functionalities to support and integrate intelligence into the framework and its services are also described. Lastly, the implementation of user preference modeling in Caravels is described.

In the fifth chapter, the developed case studies to test and analyze the framework when applied to real data and real contexts are described. Finally, in the last chapter are presented the final considerations and conclusions of the work and possible future developments.

2 State of the art

This chapter addresses the state of the art of the most relevant topics used in this thesis: intelligent communities, intelligent buildings, user preference modeling, and artificial intelligence. Besides exploring the works proposed in these domains individually, the applications of IoT devices and artificial intelligence on Intelligent Communities and Intelligent Buildings are also analyzed.

2.1 Intelligent Communities

An intelligent community can be defined as a networked socio-technical ecosystem composed of interconnected intelligent buildings and shared digital infrastructure, which collectively deploy artificial intelligence, distributed computing, and communication technologies to enhance the quality of life, sustainability, and governance of a defined population (Komninos, 2006; Prandi et al., 2020).

The concept of intelligent communities has its roots in the late 20th century, influenced by the convergence of technological advancements, increasing internet access, and the awareness of the need for sustainable urban development (Weiser, 1991). The early stages of development focused primarily on the deployment of sensors, automation and communication technologies to enable data collection and remote monitoring of infrastructures, and to improve urban services such as transportation and energy distribution (Bohn et al., 2004).

As the internet and communication technologies evolved, the concept evolved towards a citizen-centric approach, however, the emphasis was on providing digital access and information, often relying on static rule-based logic and centralized platforms (F. Y. Wang et al., 2007; Xia et al., 2015). These systems lacked autonomy and active participation from community members.

Supported by the advances in computation, artificial intelligence, and containerization technologies, the intelligent community paradigm surged as a response to these limitations, emphasizing real-time adaptability, collaborative environments and decentralized decision-making.

There is a relative difference between the concept of intelligent communities and smart communities, being the first considered an evolution of the latter (Z. Chen et al., 2022). Table 1 shows the difference between these two concepts. Unlike these smart communities, which often rely on a centralized system to achieve control and static automation, intelligent communities are characterized by their distributed, adaptive and participatory architecture, integrating several intelligent agents at different levels, such as in individual devices, buildings, or community-wide programs (Shrestha et al., 2025).

Intelligent communities strike a balance between individual goals and collective goals, through context-aware decision mechanisms and real-time resource optimization, considering evolving preferences, behaviors and needs of each member (Adel & HS Alani, 2024).

Table 1 Comparison between smart communities and intelligent communities

Dimension	Smart Communities	Intelligent Communities
Primary Focus	Monitoring; Automation; Reactive control	Learning; Adaptation; Proactive optimization
Core Technologies	IoT sensors; Embedded systems; Wireless communication	AI models; Optimization algorithms; Multi-agent systems
Decision Making	Rule-based; Deterministic	Data-driven, probabilistic, adaptive
Data Usage	Primarily for real-time control and monitoring	Data analysis for prediction, personalization, and optimization
Adaptability	Fixed Control Strategies	Context-aware; Adaptive policies

A parallelism can be traced between intelligent communities and energy communities, whereas both aim to promote community resilience and resource optimization, but the first as a stronger emphasis on human-centric design, and the latter is more focused on energy-related applications (Nishi & Nakamura, 2020).

Being data-driven, intelligent communities have the potential to promote initiatives in several domains, such as sustainability and resilience, by utilizing mechanisms to support efficient resource management, promote innovation, by creating new market opportunities for local economy, and improving quality of life, by fostering citizen participation in governance and community development (Castro & Lopes, 2022). The concept of intelligent communities recognizes citizens not only as service users or data sources but as active participants in the co-creating of an environment that prioritizes their preferences and their needs (Madrazo et al., 2025).

The development of intelligent communities is not free of challenges and potential limitations. When collecting and processing data from individuals there is a need to address privacy and security concerns, as these systems often collect sensitive information, which is vulnerable to misuse, unethical manipulation, and cybersecurity attacks (Ismagilova et al., 2020). Also, coordinating multiple stakeholders in complex ecosystems introduces demand for scalable systems that ensure fairness, and transparency, while avoiding bottlenecks, and instability (Yorgos et al., 2019). Intelligent communities support systems must also account for futureproofing, as rapid technological advancements can make solutions obsolete.

Recent literature has revealed the importance of considering the expected human participation and adaptive mechanisms, highlighting the increasing reliance on technology and data sets in understanding human behavior. Despite covering several topics within intelligent communities, the current research has several gaps. While individual systems and algorithms are well-studied, holistic frameworks that address the integration of these different components and ensure stability and interoperability between them are much less pronounced. The evaluation of long-term sustainability, considering environmental, societal, and economic impacts is another important topic to investigate, as the current solutions lack futureproofing and do not consider unexpected events. The widespread adoption of intelligent community systems remains particularly challenging due to the substantial infrastructural investment they demand and the risk of digital exclusion among residents with limited technological access or digital literacy, potentially exacerbating existing social inequalities.

2.1.1 Community Management

Intelligent communities generate a vast amount of data from various sources, creating challenges in collecting, storing, and analyzing this volume of data. Meaning that to achieve their objectives, intelligent communities need first to address big data challenges. To further facilitate this huge demand for resources to support big data analytics, cloud computing stepped in and offered an elegant and efficient solution, presenting itself as a solution to resource intensive applications (Al Nuaimi et al., 2015). The high number of sources means the data collected from intelligent communities is often heterogeneous, meaning it comes in different formats, structures, and quality levels. This heterogeneity can make it difficult to integrate and analyze data effectively. In (Rubí & de Lira Gondim, 2021a) are presented these same challenges, the lack of a common data format and sharing standard, the heterogeneity in networking and sensor technologies, and, additionally, the lack of standardized definition of environmental indicators and finally the interoperability between IoT solutions.

Data analytics, the process of collecting, cleaning, and analyzing data, is a powerful tool that can help communities identify patterns, trends, and anomalies in their data. This information can then be used to make informed decisions about community operations. For example, data on traffic patterns, public health, and general public necessities can be used to target resources and improve safety (Gracias et al., 2023). Data on community engagement can be used to identify underserved groups and develop programs to address their needs. Data on economic development can be used to attract businesses and create jobs. In addition to identifying patterns, trends, and anomalies, data analytics can also be used to predict future events. This information can be used to prepare for emergencies, such as natural disasters or public health outbreaks (V. Chang, 2021).

As communities evolve, so do the governance models that govern them. With the introduction of intelligent communities, traditional top-down models are giving way to more participatory and collaborative approaches that emphasize stakeholder engagement, transparency, and accountability (Ciasullo et al., 2020). These emerging governance models are essential for ensuring that intelligent communities are inclusive, equitable, and sustainable.

Current research remains inconclusive about the impact of smart governance practices on the livability of cities, defined as a combination of reduced social disparities, increased ecological richness, and improved economic well-being (Tomor et al., 2019).

Distributed and decentralized governance models hold the potential to address some of the shortcomings of centralized systems, particularly in terms of responsiveness, representation, and inclusivity (D. Yang et al., 2022). By distributing decision-making authority, these models empower individuals and communities to participate more actively in shaping their own governance, fostering a sense of ownership and collective responsibility, while improving the quality of life (De Guimarães et al., 2020).

While these models offer promising avenues for enhancing governance, they also present challenges related to coordination, conflict resolution, and the potential for fragmentation (T. A. Oliveira et al., 2020). Balancing decentralization with the need for effective decision-making and coordination is an ongoing area of research and experimentation.

Intelligent communities must be held accountable for their actions and decisions. This requires clear performance metrics and regular reporting to stakeholders. It also requires mechanisms for citizens to hold their governments accountable for their performance. This could include people, independent audits, and the right to information laws.

2.1.2 Resource Sharing in Communities

By leveraging digital technologies and platforms, intelligent communities can connect individuals, businesses, and organizations with idle or underutilized resources, facilitating their sharing and exchange. This approach addresses resource scarcity and optimizes resource allocation, reducing environmental impact and enhancing community resilience. To achieve this, mechanisms that enable this exchange of resources must be designed and consider the different nature of those resources, meaning that energy will not be transitioned the same way as water or even food.

As for energy utilization, major breakthroughs have been achieved, due to more interest in research and to the recent increase in decentralized energy production, which is not as greatly reflected in other areas. Solutions such as a Smart Grid can optimize electricity distribution and reduce energy consumption by using real-time data to balance supply and demand (Barone et al., 2023; Espe et al., 2018).

Beyond energy, shared infrastructure and facilities play a vital role in the efficient use of communal assets (Bousbiat et al., 2023). Intelligent systems manage access to shared resources, such as electric vehicle (EV) charging stations, coworking spaces, or on-demand fabrication tools, through context-aware scheduling and real-time reservation platforms (Hernandez et al., 2025; Rubí & de Lira Gondim, 2021b). These platforms use historical usage data, predictive modeling, and dynamic prioritization to ensure fair access and reduce idle time. By doing so, intelligent communities enhance the return on investment for high-cost infrastructure and promote equitable usage among residents (Sugandha et al., 2022).

To support inclusive governance, community-driven decision support systems are emerging as platforms for participatory resource management. These systems integrate user preferences, environmental constraints, and community-level goals into AI-based decision models (Uslu et al., 2021).

A key emerging domain is urban farming, where intelligent communities facilitate the shared production and distribution of food. Agriculture and farming revealed to be very under looked sectors. Precision agriculture can use sensors and data analytics to optimize irrigation and fertilization, leading to increased crop yields and reduced water usage (Ayoub Shaikh et al., 2022). Rooftop gardens, vertical farms, and community greenhouses can grow crops in controlled environments, and with AI systems optimizing irrigation, lighting, and harvesting schedule, such crops can use less water and land than traditional agriculture (van Delden et al., 2021).

The sharing of data and knowledge is another critical enabler of resource optimization. Intelligent communities rely on the aggregation and anonymization of data collected across various buildings to generate insights into usage patterns and system performance. This data supports distributed learning architectures that improve predictive models without violating privacy constraints (Himeur et al., 2023). Community-wide data sharing serves as a foundation for continuous improvement and collective intelligence, driving better outcomes across resource management, comfort regulation, and fault detection (Kameda et al., 2022; R. Wang et al., 2021).

In an increasingly digital and AI-driven landscape, the sharing of computational resources has become a practical consideration. Distributed computing frameworks allow surplus computational power, for example, from home servers, idle edge devices, or institutional clusters, to be shared across the community for tasks such as AI model training, simulation, or video processing (Xiao et al., 2020). This is particularly beneficial for supporting applications that require occasional high-performance computing but do not justify dedicated infrastructure per building or user.

Despite the technologies that may be developed or solutions that are created, the key element for resource sharing in any sector is the human being. As technically advanced as a solution can be, if humans are not willing to share. Cooperation seems to be natural but a difficult thing to achieve. The implementation of economic models and incentive mechanisms can motivate cooperative behaviors in shared environments. Communities are experimenting with dynamic pricing schemes, gamified feedback systems, and token-based reward structures to align individual actions with collective goals. Distributed ledger technologies, such as blockchain, may be employed to manage transactions transparently, ensuring fairness and auditability in resource exchanges.

2.1.3 Energy Community

In 2019, the EU officially recognized and established a legal framework for Energy Communities as part of the European Commission's Clean Energy Package (*DIRECTIVE (EU) 2019/944 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL*, n.d.). This framework acknowledges the rights of citizens and communities to engage directly in the energy sector.

An energy community is a type of community-driven institution that takes social control of shared energy resources through decentralization (Frieden et al., 2019). Individual consumers, producers, and prosumers in an enclosed topology can establish such space in order to build independent projects and actively contribute to a more sustainable paradigm (Gjorgievski et al., 2021).

They are characterized by their focus on collective ownership, local autonomy, and participatory governance, often functioning through dedicated microgrids, which can function alongside or independently from central utility grids (Lowitzsch et al., 2020).

The effectiveness of energy communities is further enhanced by shared energy storage systems. These systems buffer excess energy generated during peak production periods and release this energy during high-demand intervals or grid outages (Elalfy et al., 2024). AI-based controllers manage charge and discharge cycles dynamically, balancing reliability, battery health, and economic efficiency. Shared storage systems contribute to demand-side flexibility, reducing peak load impacts and increasing energy self-sufficiency (Engelhardt et al., 2022).

As local energy communities become more common, it is essential to consider the balance between individual and collective objectives. The trade-off approach (Coignard et al., 2023) provides a framework to compare the efficiency of any controller in a two-dimensional space with individual benefits versus collective constraints. In particular, this framework enables the performance evaluation of data-driven methods that lack a theoretical best-case reference. Consequently, the proposed framework plays a crucial role in ensuring that Energy Management Systems (EMS) are acceptable to end-users.

In this context, peer-to-peer (P2P) energy sharing has emerged as a platform that can facilitate the independent decision-making process of prosumers to trade their energy within a connected community (Tushar et al., 2019). In P2P sharing, prosumers can determine various parameters, such as how much energy to share, at what price, determine who to share the energy with, and when to share. It is important to note that in P2P trading allows for significant independence, a centralized controller or third party may still have some influence over the decision-making process (Tushar et al., 2021).

In EU, there are already several energy communities deployed. In (Heaslip et al., 2016) the results of preliminary fieldwork investigations of sustainable energy community development approaches in two Danish islands and one rural village in Ireland are reported. The authors' major evidence suggests that social barriers are interrelated and frequently reinforce each other, posing the most significant challenge to the deployment of energy communities. In (Giordano et al., 2019) it is proposed the development of a smart energy community, which would bring savings of 66% relative to energy costs, as the community would have a self-consumption of nearly 95%, meaning less power would have to be transported on the network.

Lastly, energy communities gather significant environmental and social benefits. By promoting the local production and consumption of renewable energy, they reduce greenhouse gas emissions and dependence on fossil fuels. Energy communities serve as an example of citizen-centric, sustainable development within the broader context of intelligent community systems.

2.2 Intelligent Buildings

An intelligent building is an adaptive and interconnected physical environment that integrates technology, namely artificial intelligence, to create advanced environments to manage and improve its operations (Clements-Croome, 1997). These operations are, usually, provided by software services, that process data acquired from sensors present in the building. Unlike, smart buildings, which usually rely on hardware for automation purposes through monitoring and controlling, intelligent buildings are characterized by their context-aware, learning and decision-making capabilities (K. Zhang, 2021). They introduce AI solutions and computing capabilities to continuously adapt to internal and external conditions, learning from historical data and environmental context, to meets its users preferences and needs (Raman et al., 2025). This means that the intelligent building are centered around its user's objectives and requirements (Dwivedi et al., 2022).

When considering that intelligent buildings are integrated in intelligent communities, is important to recognize that these buildings are not isolated entities; quite the contrary, they are interconnected nodes within a broader network, capable of communicating and collaborating with other buildings, infrastructure, and community systems to achieve both individual and shared objectives. This approach leads to a synergistic ecosystem where resources, data, and services are seamlessly exchanged to optimize community-level performance and promote sustainability (X. Chen et al., 2020; García-Monge et al., 2023a).

There are technical aspects that need to be implemented in order to build an intelligent environment within the building. At the foundational level, a sensing and controlling infrastructure is necessary, composed of distributed internet-connected devices capable of real-time data collection and actuation on environmental conditions and building subsystems

(Pathmabandu et al., 2023). A communication network and an interoperability protocol are essential to enable data exchange and communication between diverse building systems, devices, and external networks. However, having these two components is not enough, a computational layer is also required to perform the data processing and provide results. important to create edge and cloud computing, which provides distributed data processing and storage capabilities, enabling real-time analytics (Schmidt et al., 1987).

For computation, there are two possibilities, either the building integrates cloud computing or local computing (Anawar et al., 2018). At its core, cloud computing refers to the delivery of computing services, these including servers, storage, databases, networking, software, analytics, and intelligence over the internet (Ghazy et al., 2023). Edge computing refers to a distributed computing framework that brings data storage and processing closer to the location where it is needed, improving response times and saving bandwidth (Marcham, 2021). Typically, cloud computing provides more computational resources and power when compared to edge, which when integrating AI models can be more resource-constrained (Caiazza et al., 2021). To mitigate this problem it is possible to support both solutions, facilitating local decision-making while levering centralized resources for more complex processing or more computation demanding tasks.

As already mentioned, central to intelligent behavior is the implementation of AI-based control and optimization, where learning algorithms perform predictive modeling, detect anomalies, and dynamically schedule operations (Merabet et al., 2021). In order to allow dynamic adaptation, AI models have to integrate user preferences, translating occupants preferences, presence patterns, and feedback. This is key for aligining building behavior with human-centric objectives (Nabizadeh Rafsanjani & Nabizadeh, 2023).

Current research trends are advancing rapidly across multiple domains. Adaptive control, through reinforcement learning, such as in (Ribeiro et al., 2024), is being used to optimize HVAC operations, reducing energy usage while maintaining thermal confort. To address privacy and latency concerns, federated and edge learning paradigms are being explored to enable predictive maintenance, while preserving data locality. In (Cespedes-Cubides & Jradi, 2024) was demonstrated the use of digital twins to simulate building performance, and to evaluate control policies within commercial buildings.

In summary, intelligent buildings are not merely technological upgrades to the built environment, but key enablers of sustainable, adaptive, and collaborative ecosystems.

2.2.1 IoT on Smart Buildings

Internet of Things constitutes the sensory and control backbone of intelligent buildings, acting as a physical-digital interface between the building and the integrated systems (Verma et al., 2019). By deploying a network of distributed sensors and actuators, IoT enables continuous monitoring and control of building systems. These networks of interconnected devices collect data across spatial and temporal scales, enabling autonomous adaptation of building operations and facilitating communication with other entities (García-Monge et al., 2023b).

A diverse range of sensors can be deployed throughout intelligent buildings, such as temperature, humidity, CO_2 concentration, occupancy, luminance, motion, and air quality sensors. These devices are distributed according to architectural topology and operational needs, with placement strategies that maximize coverage while minimizing redundancy and energy consumption. For instance, occupancy can be inferred more accurately through the combination of motion and CO_2 data, whereas precision temperature control in multi-zoned HVAC systems requires granular thermal sensing across rooms and corridors (Ding et al., 2021). Additionally, with the inclusion of newer technologies, such as wearables, the amount of information and solutions available increases exponentially (John Dian et al., 2020; Singh et al., 2020).

In parallel, actuators provide a mechanical interface for implementing control decisions. Examples include lighting dimmers, blinds, or even infrared emitters. When integrated with AI models, actuators enable closed-loop feedback systems that adjust dynamically to the changing conditions. The coordination between sensors, actuators, and the AI models compose an ecosystem capable of responding to multiple building objectives.

Reliable and low-latency communication protocols are essential to the effective functioning of IoT systems (Lawal & Rafsanjani, 2022). Lightweight standards, such as Message Queuing Telemetry Transport (MQTT), ZigBee, LoRaWAN, and Bluetooth Low Energy (BLE), are commonly employed to support the energy-efficient transmission of data across constrained devices. However, despite being able to transmit data, a critical challenge is to guarantee interoperability between heterogeneous devices and systems (Noaman et al., 2022). Devices often originate from different manufacturers, adhere to varying data formats, and operate under incompatible communication protocols, resulting in fragmented and semantically inconsistent data streams. This heterogeneity impedes the seamless aggregation and interpretation of sensor data necessary for higher-level reasoning and decision-making (Qiu et al., 2018). To address this issue, middleware solutions, such as Matter[8], or open Application Programming Interface (API), as provided by Home Assistant[9], enable the standardization of data and cross-device compatibility.

The design of IoT architectures in intelligent buildings prioritizes modularity and scalability, allowing for incremental expansion and adaptation to evolving requirements (Ibaseta et al., 2021b). Modular sensor and actuator nodes enable rapid deployment, configuration, or maintenance without disrupting existing infrastructure. While scalable communication protocols and data management systems allow for expansion across building types and sizes, whether it is a single residential home or a large office building (Djehaiche et al., 2023).

[8] https://handbook.buildwithmatter.com/
[9] https://www.home-assistant.io/

Nevertheless, the exponential adhesion of IoT devices raises major concerns in security and privacy. Given the attack surface exposed by wireless connectivity and distributed endpoints, intelligent buildings must employ cybersecurity measures, including encrypted communication channels, authentication protocols, and intrusion detection systems (Schiller et al., 2022). Additionally, with the collection of sensitive data, there are certain questions that arise, such as where is the data stored, who has access to the data, and how it is used (Gupta et al., 2022).

2.2.2 Resource Optimization in Buildings

One of the central functions of intelligent buildings is the optimization of resource utilization, while maintaining or improving user comfort and operational reliability. Traditionally, building resource management has been carried out through scheduled operations or rule-based controls, which are inherently limited in their responsiveness to dynamic environmental conditions and occupant behavior (Freund & Schmitz, 2021). In contrast, intelligent buildings employ AI-driven models to balance multiple, and often conflicting objectives. These models can operate in real-time enabling optimizations, such as minimizing energy consumption, reducing peak load demand, and maintaining thermal, visual, and acoustic comfort (Esrafilian-Najafabadi & Haghighat, 2021). Such actions are conducted under constraints imposed by the physical infrastructure, user requirements, or external signals, such as electricity prices or weather forecasts.

The usage of AI models with predictive capabilities is fundamental to this problem, enabling systems to anticipate future states and preemptively adjust control parameters, bases on learned models (Olawade et al., 2024). For instance, AI models can learn from historical occupancy patterns, weather forecasts, and building thermal characteristics to predict future energy demand and proactively optimize HVAC and lighting settings (S. Lee & Karava, 2020).

Controlling discrete operations can be achieved through dynamic scheduling algorithms. These algorithms allocate and schedule resources or activities, such as appliance use, lighting cycles, or battery charging/discharging, based on predictions, constraints or objectives, usually in a way that reduces costs, aligns with renewable energy generation, or avoids peak loads (Q. Yang & Wang, 2021). Techniques such as mixed-integer programming, constraint programming, or more recently, deep reinforcement learning have shown promise in handling complex, non-linear scheduling problems in real-time, especially under uncertainty (O'Malley et al., 2023). Nevertheless, the deployment of optimization models requires a holistic understanding of building physics and operations, coupled with high-quality, real-time data streams from diverse sensors and systems.

Finally, intelligent buildings are beginning to coordinate their operations at the community scale. In this context, building exchange information with each other to optimize collective resource utilization and resilience, while preserving autonomy at the individual level (Nik & Moazami, 2021). This coordination relies on communication infrastructure, and, for example, multi-agent systems that allow for scalable and distributed decision-making, which seek to explore both global and individual goals.

2.3 User Preference Modeling

Modeling preferences is not a trivial challenge, this is due to the inherent psychological complexity of humans (Schürmann & Beckerle, 2020). It is important to underline the organic and situated nature of human preferences. Rather than being universal or static, preferences are context-dependent, dynamic, subject to cognitive and emotional influences, and can be constructed at the moment of decision (Terenzi et al., 2021). Cognitive psychology has shown that individuals frequently lack introspective awareness of their own preferences and may struggle to articulate them without explicit prompts (Langford et al., 2020).

Preferences are prone to biases, framing effects, and social desirability considerations, leading to inconsistencies between stated and revealed preferences (Beshears et al., 2008). Additionally, preferences can be inconsistent across time or domains, being influenced by current emotional states, social context, or previous choices, such as the endowment effect or choice-induced preference change (Dietrich & List, 2013). People have the tendency to favor options they have previously chosen, even if those options were initially arbitrary. One well known example of this bias is placebo effect, which is a phenomenon where a patient experiences improvement in their condition after receiving an inactive treatment (Meppelink et al., 2019). This improvement is due to the patient's belief in the treatment and their expectation of feeling better, rather than the actual properties of the treatment (Pardo-Cabello et al., 2022). This anthropological behavior reflects how malleable, endogenous, and sometimes abstract, human preferences can be. These psychological dynamics complicate any attempt to formalize preferences in static or rational terms.

Conventional preference elicitation techniques fall short of capturing the complexities and nuances of human cognition (Atas et al., 2021). Traditional methods often rely on direct questioning or surveys, which are susceptible to biases, demand effects, and introspective limitations, failing to capture these nuances and contradictory nature of what user actually want (Cheng et al., 2020). For example, a user might state that he prefers natural lighting but closes the blinds when to avoid glare or thermal discomfort. Additionally, there is great difficulty in converting subjective preferences into quantifiable or measurable metrics. Therefore, effective preference modeling both explicit and implicit inputs should be used, leveraging observations, sensor data, and behavioral patterns to infer underlying preferences, while adapting to environmental and contextual cues (Pigozzi et al., 2015).

From a representative perspective, user preferences must be embedded in a coherent knowledge representation framework that allows interactions with AI models. Symbolic approaches, such as rule-based, ontological hierarchies, and logic-based representations, provide transparency and traceability, making them useful for systems that requires explainability and verification (Auda et al., 2018). In contrast, connectionist or subsymbolic approaches leverage neural networks or hidden representations to provide greater scalability and adaptability by learning from data, while sacrificing interpretability (Guo & Wang, 2021). Hybrid approaches try to merge the strengths of both, integrating structured logic with learned embeddings, such as neuro-symbolic approaches or graph-based approaches (Spillo et al., 2024).

Different approaches have been proposed to represent users in intelligent systems. Basic models rely on static user profiles, where preferences are stored as fixed attributes (Poo et al., 2003). While easy to implement they do not capture the dynamic nature of preferences. To this end, context-aware models incorporate environmental information to provide flexibility. Behavioral models go even further, using historical interactions and sensor data to infer hidden preferences and detect routine patterns (Kanoje et al., 2015). More sophisticated approaches utilize cognitive and affective dimensions, to capture aspects like emotion, personality, and intent (Gao et al., 2010; Ying et al., 2018).

When facing a multi-user setting, models also need to account for group dynamics, interpersonal preferences, and conflict resolution strategies. Depending on the scenario, different negotiation or priority mechanisms may be implemented, as there is not a strategy that suits all situations. For example, in a residential setting the system might differentiate residents from visitors, giving more priority to residents, while in an office setting a strategy that balances preferences with other factors, such as energy consumption, is preferred.

User preference modeling refers to the digital or computational representation of individual or group preferences, behaviors, or needs (R. Wang et al., 2022). User preference modeling is critical for the creation of adaptive, personalized and responsive environments. At its core, this process seeks to capture the nuances of human behavior and translate them into actionable insights for building automation and control systems (Zhu et al., 2021). The ability to model and act upon user preferences distinguishes intelligent systems from merely automated or reactive ones, supporting higher levels of personalization and acceptability (Mallik & Gangopadhyay, 2023). Potentially, the integration of such systems can minimize interactions with a system, by predicting the required setting before the user even needs them.

The design of user preference models involves accounting for the dynamic and contextual nature of human behavior. Preferences may vary on short and long term periods, being impacted by several factors as time of day, season, location, current or previous activities, and social context. Context-aware systems embed these variables into the decision process, reacting to the real-time conditions and user inputs to deliver the most appropriate response. Also, preferences may be reflected due to internal or psychological expressions, such as stress, emotions, mood or fatigue (Kielar & Borrmann, 2016). These biometric or affective states, when properly measured and interpreted, can provide a deeper insight into user needs and comfort levels. For example, users might prefer a certain genre of music when facing different situations.

Despite notable advancements in user modeling, significant challenges remain, particularly in adapting to the inherent variability and ambiguity in human behavior. The resolution of multi-user conflicts in a fair and privacy-preserving manners, and the integration of distributed architectures that span multiple buildings (Aljawarneh et al., 2021). There are major considerations regarding data sovereignty, user consent, and algorithmic transparency when implementing such models (Eke et al., 2019). Advanced preference modeling requires multi domain research, bridging disciplines such as artificial intelligence, cognitive science, and psychology.

2.3.1 User Interaction with Digital Systems

Historically, digital systems responded only to explicit user commands, requiring direct input for every action. However, modern intelligent systems are increasingly designed to anticipate user needs and act proactively (T. Jiang et al., 2024). Through historical data analysis, these systems can recommend, automate routine decisions, or initiate actions without direct prompts. This capability enhances efficiency and convenience, yet it also introduces considerations regarding transparency and control.

There are multiple methods through which users can interact with intelligent systems, from smartphones and wearables to building systems and voice-controlled assistants (Norval & Singh, 2024). This distributed interaction paradigm necessitates interfaces that are consistent and intuitive across multiple devices, ensuring a seamless user experience. Additionally, selecting the right interaction modality is critical to this seamless experience, so the integration of context variables becomes imperative. The proliferation of multi-modal sensing technologies, including motion detectors, voice recognition, and biometric systems, enables deeper situational understanding and more natural user-system interactions (Cosentino & Giannakos, 2023).

As digital systems grow more autonomous and pervasive, ensuring transparency, user control, and trust becomes increasingly essential. Users must understand how and why decisions are made on their behalf. To that end, explainable interfaces, adjustable automation settings, and accessible feedback channels are necessary to maintain informed participation and promote user confidence (Zieglmeier & Pretschner, 2021). Trust, which is a critical factor in long-term engagement, is cultivated not only through functional accuracy but also through ethical design principles, such as privacy protection, consent management, and accountability in decision-making.

2.3.2 Multi-user Environments

Managing interactions in multi-user environments introduces a distinct set of challenges. Unlike individual environments where the system focuses solely on a single user’s preferences, multi-user settings mean that the system must account for the coexistence of multiple users, each having their own set of preferences, which may conflict (Furszyfer Del Rio, 2022). Ensuring fairness, personalized experience while maintaining satisfaction requires the system’s capability to recognize individual users and differentiate among them.

To mediate between multiple user preferences, systems must employ preference aggregation and conflict resolution strategies. The system can have several approaches, and each approach might have different trade-offs concerning fairness, efficiency, and user satisfactions (Shin et al., 2010). Several approaches are found in the literature, such as weighted averaging, which can be based on usage history or user role, prioritization hierarchies, voting mechanisms, computational negotiation models, or rock-paper-scissors (Masthoff, 2004). Depending on the scenario, different strategies may be implemented, as there is not one that suits all situations. For example, in a residential setting the system might differentiate residents from visitors, giving more priority to residents, while in an office setting a strategy that balances preferences with other factors, such as energy consumption, is preferred (Mary Reena et al., 2015). Moreover, systems can leverage historical group interaction data to learn collective preferences over time, improving the relevance and acceptability of aggregated responses.

A distinction must be made between private and shared spaces within a defined environment. In private zones such as bedrooms or personal offices, user modeling can be deeply individualized (Von Frankenberg, 2021). In contrast, shared environments like living rooms, meeting rooms, or kitchens require more adaptive models that balance the needs of all occupants (P. Oliveira et al., 2018). Intelligent systems must be capable of dynamically transitioning between individualized and collective operational modes, based on real-time context and occupancy.

When scaling a system capable of handling multi-user environments, maintaining accurate and responsive user models becomes computationally demanding. The complexity increases with the number of users, diversity of preferences, and dynamism of interactions. Finally, the user interface must evolve to support collaborative and socially aware interaction. Interfaces should facilitate multi-user configurations, present aggregated system states transparently, and allow for collective decision-making where applicable. In some cases, systems may act as mediators, guiding users through negotiation processes or presenting trade-offs to support consensus. These capabilities not only enhance usability but also foster trust and acceptance in shared environments.

2.4 Artificial Intelligence

AI refers to a field of computer science dedicated to the development of computational systems capable of performing tasks that traditionally require human intelligence. These tasks often encompass perception, learning, reasoning, problem-solving, decision-making, and language understanding (Bibri et al., 2024; Hopgood, 2005a). The goal of AI is to create algorithms and models that can interpret complex data, draws inferences, adapt to new situations, and act in ways that achieve specific goals (Y. Jiang et al., 2022). The applicability of AI has become engraved in many sectors, such as industry, finance, healthcare, automation, and beyond.

Development of AI draws foundational concepts and insights from mathematics, psychology, neuroscience, linguistics, and cybernetics (C. Zhang & Lu, 2021). This interdisciplinary nature of AI leads the design of such systems to mimic cognitive processes and biological mechanisms, bridging the gap between artificial and natural intelligence (Zador et al., 2023). The synergy between these two domains has fueled major breakthroughs in areas such as optimizations, deep learning, and natural language processing.

Artificial Intelligence is not a recent field by any means; since its conceptualization in the mid-20th century several branches developed, each specialized to address particular challenges and objectives (Howard, 2019). Figure 1 shows a high level overview of such branches.

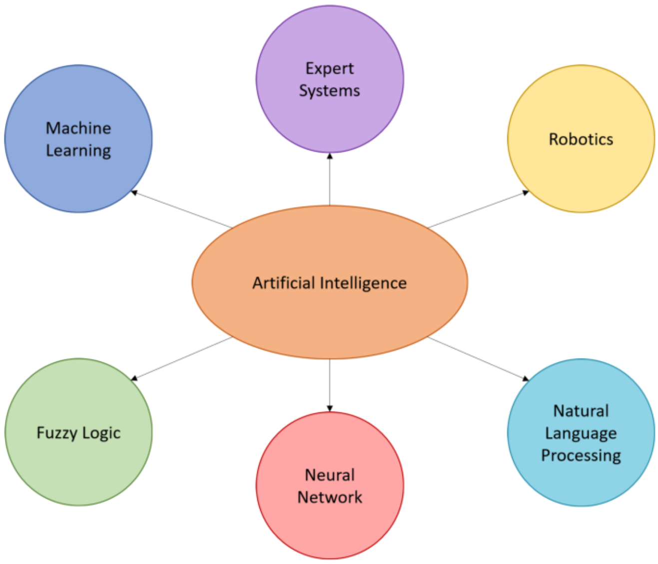

Figure 1 Branches of Artificial Intelligence. Adapted from (Butt et al., 2022)

From these branches, we can focus on three essential paradigms: symbolic methods, statistical learning, and biologically inspired models. Symbolic AI employs explicit rules and representations using logic and knowledge-based systems (Hopgood, 2005b). Statistical learning, which is known for its machine learning techniques, relies on probabilistic models to learn patterns and make predictions or classifications (Maruyama, 2022). Biologically inspired models, such as neural networks, evolutionary algorithms, or swarm intelligence, are grounded on principles observed in natural occurrences or organisms (Fister et al., 2013).

More recently, with the increasing availability of data and computing power, Machine Learning (ML) has been receiving significant attention in the literature. ML algorithms provide the advantage of improving their performance automatically as they receive new data, eliminating the need for explicit programming (Shalev-Shwartz & Ben-David, 2014). However, there are some limitations to their use. If an ML algorithm is not properly parameterized, it can lead to overfitting. This occurs when the model becomes too tailored to the training data, making it ineffective at making predictions outside the range of values it was trained on (Montesinos López et al., 2022). Under the machine learning branch, it is possible to identify three main categories: supervised learning, unsupervised learning, and reinforcement learning. Supervised learning learns from labeled data to make predictions, unsupervised learning finds hidden patterns in unlabeled data, and reinforcement learning learns optimal actions through trial and error by interacting with an environment and receiving feedback (Sun et al., 2019).

The integration of AI has become instrumental in revolutionizing urban living. A diverse array of AI algorithms contributes to the advancement of these intelligent ecosystems. ML algorithms, including supervised and unsupervised learning, find application in predictive analytics for forecasting in diverse domains and management/optimization (Aguilar et al., 2021). Reinforcement learning algorithms have notable potential optimize decision-making in dynamic systems, such as energy grids and traffic management, for resource efficiency (Perera & Kamalaruban, 2021; Yau et al., 2017).

The applications of AI in intelligent communities are multifaceted and transformative. In traffic management, AI facilitates predictive modeling, intelligent traffic light control, and dynamic route planning to alleviate congestion in urban areas (Yau et al., 2017). Healthcare benefits from AI through predictive analytics for disease outbreak, preventing pandemics such as COVID-19 (Ardabili et al., 2020). For public safety, more specifically in learning disaster and public health emergency context, development and use of AI have increased over the past recent years, leading towards a swift and efficient response to emergencies (Lu et al., 2022). Furthermore, AI represents a major role in environmental protection and sustainability not in the sense in how it will help society reduce resource consumption but rather how it facilitates and fosters environmental governance (Taghikhah et al., 2022).

Looking toward the future, several directions emerge for the evolution of AI in intelligent communities. Human-AI collaboration is anticipated to become more pronounced, fostering a symbiotic relationship where AI systems and human residents work collaboratively for more effective decision-making (Park et al., 2019). Edge computing for AI will likely gain prominence, enabling real-time processing and reduced latency (Z. Chang et al., 2021). Additionally, interoperability is crucial, ensuring seamless integration of diverse AI applications and the development of holistic intelligent community ecosystems where different systems can work together efficiently.

2.4.1 Supervised and Unsupervised Learning

Supervised learning is a key area of machine learning within AI that focuses on learning a mapping function from input data to known output labels (El Naqa & Murphy, 2015). This approach relies heavily on labeled datasets, where each input is paired with a corresponding correct output. This setup enables the model to learn from past examples and make accurate predictions on new, unseen data. The main goal is to minimize a loss function, which measures the difference between the predicted outputs and the actual outputs, thereby refining the model iteratively during training.

Common tasks in supervised learning include classification, which involves assigning inputs to discrete categories, and regression, which involves predicting continuous numerical values (Kourou et al., 2015).In the current scope, supervised learning is particularly useful in scenarios where there is a need to predict or classify outcomes based on historical data, such as predicting energy consumption patterns of buildings, classifying occupancy levels, or forecasting equipment failure based on sensor data (Fan et al., 2021).

A wide range of models are employed in supervised learning, from interpretable linear classifiers and decision trees to more complex architectures like support vector machines and ensemble methods such as random forests and gradient boosting (Zheng et al., 2020). In recent years, deep learning models, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformers, have shown remarkable performance across various tasks by leveraging large-scale data and computational power (Lecun et al., 2015).

However, training supervised learning models requires a significant amount of labeled data, which can be both time-consuming and costly to obtain. In contrast, rather than learning from predefined input-output pairs, unsupervised learning aims to model the underlying distribution of the data or identify meaningful groups based on proximity (Y. Zhang, Chen, et al., 2023). This approach can be used for exploratory data analysis, data compression, and situations where labeled data is scarce or unavailable. Figure 2 illustrates a comparison between these two methods.

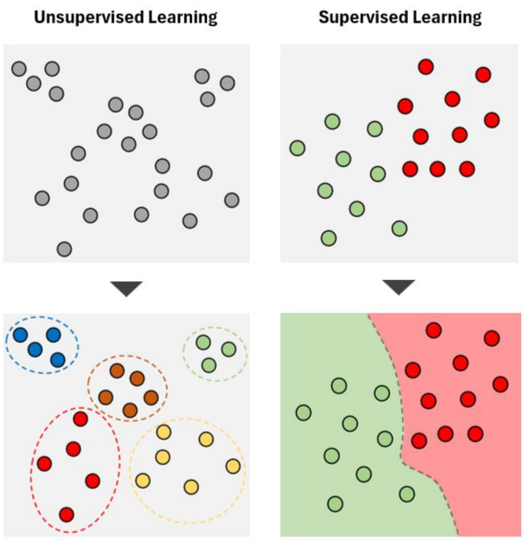

Figure 2 Comparison between supervised and unsupervised learning. Adapted from (Morimoto & Ponton, 2021)

Tasks within unsupervised learning include clustering, dimensionality reduction, anomaly detection, and representation learning. Clustering algorithms, such as K-means, hierarchical clustering, and Gaussian mixture models, aim to group similar data points based on distance or density metrics (Lam & Wunsch, 2014). Techniques like swarm intelligence and Hebbian learning draw on principles from neuroscience and collective behavior, enabling adaptive and decentralized learning systems (Marzband et al., 2014). In the current scope, unsupervised learning can be used for recommendation services, or situations where there is no prior knowledge.

2.4.2 Reinforcement Learning

Reinforcement Learning (RL) is a type of learning in which an agent/learner interacts with their environment through trial and error (Gronauer & Diepold, 2022). In contrast to other machine learning techniques, the agent learns to make decisions by interacting with an environment and receiving feedback in the form of rewards or penalties. Throughout these interactions, the agent learns to find the best strategy, which in RL is called policy, that maximizes cumulative rewards.

The theoretical foundation of RL is based on Markov Decision Processes (MDPs), which formalize the interaction between an agent and its environment through states, actions, transition probabilities, and reward functions (Hubbs et al., 2020). Various algorithms have been developed to address RL problems, which can be categorized into value-based and policy-based approaches (Liang et al., 2025).

Value-based methods focus on estimating the value function, which predicts the expected cumulative reward, also known as return, from a given state or state-action pair, assuming the agent follows a specific policy. Value-based methods, such as Q-learning, SARSA, or Deep Q-Networks, are typically more sample efficient and work best in discrete and low-dimensional action spaces (Watkins & Dayan, 1992).

On the other hand, policy-based methods directly learn the policy, which maps states to action probabilities, without relying on a value function to make decisions. These methods, such as Policy Gradient, Proximal Policy Optimization, or even hybrid methods, which combine elements from both value and policy-based methods, such as Actor Critic, are better suited for high-dimensional or continuous action spaces but less sample-efficient (Bhatnagar et al., 2009).

Additionally, RL can operate in online or offline mode. In online mode, the agent learns by interacting with the environment in real time. It observes the consequences of its actions, updates its policy or value estimates, and then uses the new policy to decide its next action (Figueiredo Prudencio et al., 2024). In offline RL, the agent learns from a fixed dataset of past interactions, collected by another agent, human, or logged system. It does not interact with the environment during training.

Human-inspired techniques, such as imitation learning and inverse reinforcement learning, allow agents to learn from demonstrations, reducing the sample complexity (Nozari et al., 2022). Hierarchical reinforcement learning introduces subgoals and temporal abstraction, facilitating more efficient learning in complex environments (Le et al., 2018).

Despite its potential, RL faces limitations. A significant one is sample inefficiency, as learning optimal policies often requires millions of interactions. Reward sparsity and delayed feedback further complicate learning, especially in tasks where the correct sequence of actions is not immediately rewarded (Lyu et al., 2024). Safety and stability during exploration remain open challenges, particularly in real-world applications where poor decisions can have costly or dangerous consequences.

3 Methods and Materials

This chapter describes the methods and algorithms used in the development of the platform, as well as the tools used both to develop and test it. Besides, this chapter presents the technological and social challenges related to the domain of development.

3.1 Materials and Tools

This section outlines the materials and tools utilized in the development of the proposed system. It covers key software and hardware technologies, development frameworks employed, and the data collected.

3.1.1 Software Infrastructure Enabling Technologies

In the context of distributed applications, robust and scalable software infrastructure is essential. Intelligent systems today are often deployed across heterogeneous environments, requiring modularity, high availability, rapid scalability, and efficient resource usage. This section outlines the foundational technologies that enable such infrastructure, namely microservice architecture, containers, container orchestration, and Virtual Private Networks (VPNs). These technologies are especially relevant for applications that rely on dynamic service composition, real-time data processing, and secure cloud-native deployment environments.

Microservice architecture represents a change from traditional monolithic system design, emphasizing the division of applications into smaller, self-sufficient services. Each microservice encapsulates a distinct business function and operates as an independent deployment unit, runtime environment, and lifecycle. This modularization allows development teams to iterate and deploy features rapidly, isolate faults, and scale individual components based on demand rather than duplicating the entire system.

Microservices communicate primarily through lightweight protocols such as Hypertext Transfer Protocol (HTTP)/Representational State Transfer (REST) or message queues, enabling loose coupling and language-agnostic implementation. This flexibility is particularly beneficial in AI ecosystems, where different components, such as model inference engines, data ingestion pipelines, and logging services, may be developed in different languages and require independent scaling. While most models are written in Python, the interpreted nature of this language presents major disadvantages, namely in situations where performance is a critical factor. In such cases, compiled languages may offer better execution speed and resource efficiency. With a microservice architecture, these compiled languages can be utilized in performance-sensitive components without creating dependencies across the entire system. However, microservices also introduce disadvantages such as increased complexity in managing distributed systems, and overhead and latency in service communication. To address the distributed characteristics of the system, the system serves itself with three key technologies: containers, container orchestration, and VPN.

Containers are a fundamental enabling technology for deploying microservices. A container packages an application along with all its dependencies, libraries, and environment configurations into a single, portable unit. Table 2 presents a comparison between container and virtual machines (VMs).

Table 2 Container and Virtual Machine comparison

Feature	Container	Virtual Machine
Virtualization Type	OS-level (shares host OS kernel)	Hardware
Startup Time	Seconds	Minutes
Resource Usage	Low (no full OS per container)	High (each VM includes full OS)
Isolation	Process-level	Full OS-level
Portability	High	Moderate
Performance Overhead	Minimal	Higher due to OS emulation
Use Case Suitability	Microservices; DevOps; CI/CD; AI	Legacy apps; Full OS simulation

This architectural efficiency results in lightweight, fast-starting environments with reduced memory and Central Processing Unit (CPU) overhead, while also ensuring consistency across environments. The portability of containers allows AI models trained in one environment to be seamlessly deployed in another, supporting reproducibility and collaborative development in multi-machine infrastructures. The proposed solution uses Docker[10] to provide a platform to package, distribute, and manage applications within containers. Unlike virtual machines, which virtualize at the hardware level and require separate operating systems, Docker uses Operating System (OS) level virtualization, sharing the host kernel while isolating applications at the process level, as seen in Figure 3.

[10] https://docs.docker.com/

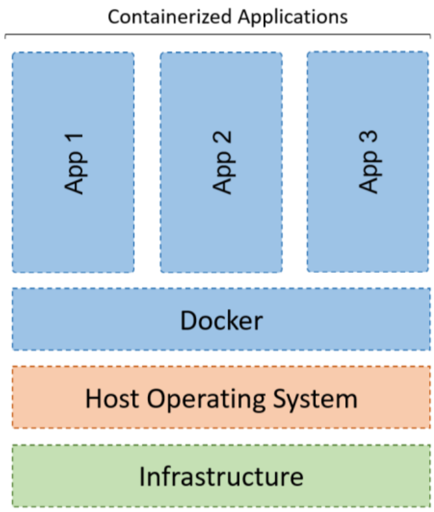

Figure 3 Docker architecture

While containers simplify application packaging and deployment, managing hundreds or thousands of them across multiple hosts introduces operational complexity. Container orchestration platforms, such as Kubernetes[11] which is open source, address this challenge by automating the deployment, scheduling, scaling, and managing containerized applications. A Kubernetes system is complex and has a lot of concept definitions that, in order to understand how it all comes together, need to be broken down.

Kubernetes follows a cluster architecture, as exemplified in Figure 4, that consists of a control plane plus a set of worker machines, called nodes, that run containerized applications, called Pods. Every cluster needs at least one worker node in order to run Pods, which are the components of the application workload. The control plane manages the worker nodes and the Pods in the cluster. For simplicity, this dissertation will not delve into each component of the control plane or the node, as they are abstracted. Nonetheless, it will provide a definition of the more general concepts:

- Cluster: a set of nodes that work together to run containerized applications in a scalable, resilient, and automated manner. It consists of worker nodes that run the application workloads and control plane components that manage the cluster's state and operations;
- Node: a physical or virtual machine that runs containerized applications. Serves as the execution environment for Pods;
- Pods: a group of one or more containers, with shared storage and network resources, and a specification for how to run the containers. A Pod acts as a wrapper around a container.

These orchestration platforms provide a control plane that manages containerized applications, ensuring high availability, fault tolerance, and efficient resource allocation. They monitor the health of services, restart failed containers, perform rolling updates, and support declarative

[11] https://kubernetes.io/docs/

configurations through manifest files. This improves the resilience and scaling of the system as well as facilitates the management of the individual configuration of each smart building.

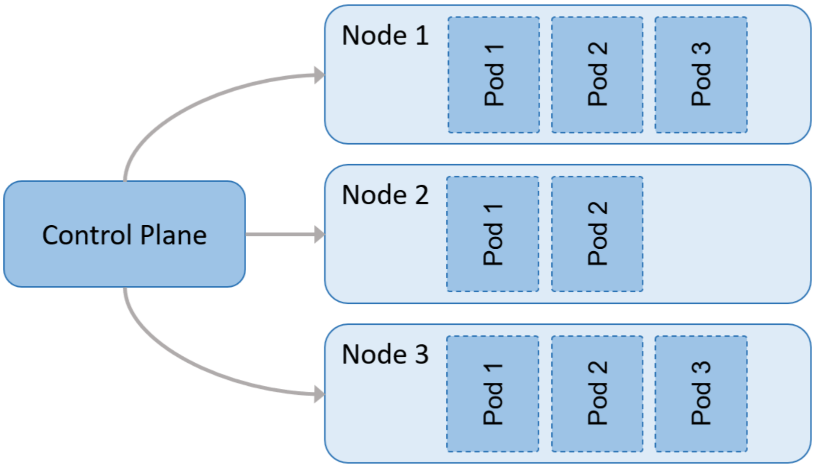

Figure 4 Kubernetes architecture

In distributed systems, services and users, which may be spread across different locations and networks, need to communicate between them. To enable secure and stable communication, the system uses a VPN which establishes a digital connection between two machines, creating an encrypted point-to-point tunnel that channels over public and shared networks, avoiding exposing firewall ports to the external internet.

VPNs are commonly used to connect remote developers to internal infrastructure, protect sensitive APIs, and establish secure links between microservices deployed across public cloud and on-premise environments. In intelligent communities, VPNs are fundamental to protect inter-building communications from unauthorized access or interception.

There are several solutions to implement VPNs, Tailscale[12] was the one chosen. Instead of implementing a “hub and spoke” architecture, where each client device connects to a central VPN gateway. Using Headscale[13], the open-source version of Tailscale, the system implements a distributed communication architecture that closely resembles the already mentioned architecture designed for this project.

3.1.2 Hardware Infrastructure Enabling Technologies

The successful implementation of intelligent buildings and communities relies not only on software but also on a diverse hardware infrastructure. As AI systems increasingly proliferate in real-world environments, physical components must support the acquisition, processing, and actuation of data in a distributed and resource-constrained settings. This section outlines the key hardware technologies that enable such systems, ranging from IoT devices and edge

[12] https://tailscale.com/
[13] https://headscale.net/

computing platforms to centralized servers and energy storage systems. These components facilitate real-time perception, local intelligence, centralized analytics, and autonomous environmental control. To support data acquisition and to enable the real world iterability capabilities of the system, several machines and IoT devices equipped with sensors and actuators will be deployed. Table 3 provides a detailed list of each device, along with the sensors and actuators it includes, as well as their specific function. Each device has a specific purpose. A single-board computer (SBC), due to its high computational capability and energy efficiency, stands out among the others. It can host services, run lightweight AI models, and connect with external equipment.

Table 3 Edge and IoT devices

Name	Type	Communication	Sensors	Use of IoT
Broadlink RM4	Universal Remote	Infrared; Radio Frequencies	None	Establishes a bridge to control HVAC
ESP32-Humidity	Development Board	Wi-Fi; Bluetooth	Humidity	Monitor humidity in a given room
ESP32-Infrared	Development Board	Wi-Fi; Bluetooth; Infrared	Infrared	Control HVAC devices and/or detect signals sent from a remote controller
ESP32-Movement Detector	Development Board	Wi-Fi; Bluetooth	Motion Sensor	Detect human presence in a given space
ESP32-Temperature	Development Board	Wi-Fi; Bluetooth	Temperature	Monitor temperature in a given room
Raspberry Pi 4	SBC	Wi-Fi; Bluetooth	None	Software execution
Shelly EM	Energy Monitor	Wi-Fi	Energy	Monitor energy consumption or generation in a building
Shelly Plus 1PM	Smart Relay	Wi-Fi; Bluetooth	Energy	Control and measure the consumption of a device
Shelly Pro EM-50	Energy Monitor	LAN; Wi-Fi; Bluetooth	Energy	Monitor energy consumption and/or generation in a building
Sonoff S26	Smart Plug	Wi-Fi	None	Controls On/Off operation of a socket
TPLink HS100	Smart Plug	Wi-Fi	None	Controls On/Off operation of a socket
TPLink HS110	Smart Plug	Wi-Fi	Energy	Real-time power consumption; Controls On/Off operation of a socket

While edge devices manage localized processing, centralized servers remain essential for computationally intensive tasks such as training deep learning models, managing large-scale inference services, and orchestrating system-wide services.

Smart appliances, such as connected thermostats, lighting systems, refrigerators, and washing machines, are becoming increasingly common in intelligent buildings and are used when available. These devices are equipped with embedded computing capabilities and wireless communication interfaces, allowing them to be monitored and controlled remotely.

Finally, energy production and storage is an essential component of intelligent infrastructure, particularly to enable self-sufficiency and resilience. As the demand for sustainable and uninterrupted energy grows, especially in environments relying on AI-driven automation and constant connectivity, the integration of renewable energy sources such as solar panels along with energy storage systems becomes mandatory.

3.1.3 Frameworks and Libraries

The implementation of intelligent systems, particularly in dynamic environments such as buildings and communities, relies on versatile software tools. The Python ecosystem provides a rich set of libraries and frameworks that support every stage of any software and AI development lifecycle, from data preprocessing to model training, evaluation, deployment, and real-time interaction. This section presents the main frameworks and supporting libraries used in the development of the proposed system, emphasizing tools selected for their maturity, community support, and compatibility with edge and cloud-based deployment scenarios.

TensorFlow, developed by Google, is a comprehensive open-source framework for large-scale machine learning and deep learning applications. It offers both high-level APIs, such as Keras, for prototyping and low-level operations for fine-grained model control. It also supports deployment across mobile and embedded devices through platforms like Tensorflow Lite. PyTorch, developed by Meta, is favored in research and prototyping due to its dynamic computation graph and intuitive Python interface. PyTorch provides excellent support for Graphics Processing Unit (GPU) acceleration and is widely used for tasks involving computer vision, natural language processing, and reinforcement learning. Scikit-learn complements the above deep learning frameworks by providing a reliable and easy-to-use toolkit for classical machine learning algorithms. It includes implementations of support vector machines, decision trees, ensemble methods, clustering algorithms, and evaluation metrics.

Efficient data handling is crucial in any AI pipeline, especially when working with high-frequency sensor data and large volumes of event logs. Two foundational Python libraries are used in this context: NumPy and Pandas. NumPy provides fast and efficient array manipulation capabilities and is the underlying structure for most numerical computations in Python-based machine learning. It supports multi-dimensional arrays, mathematical operations, and linear algebra functions. Pandas builds on NumPy to offer high-level data structures such as DataFrames, which can manipulate tabular and time-series data. Its functionality includes filtering, merging, resampling, and statistical summaries. These capabilities are useful when cleaning and organizing heterogeneous data collected from IoT sensors in intelligent buildings.

In the visualization category, Matplotlib, Streamlit, and Vue.js are employed to support data analysis and presentation of outputs to end-users. Matplotlib is a plotting library for Python, that is used to create static graphs when evaluating model performance, analyzing data distribution, and observing models metrics. Streamlit provides a lightweight and fast to deploy web-based dashboard and front-end application directly in Python. It is especially useful for creating quick front-end prototypes when developing proofs of concept. In contrast, Vue.js is a progressive JavaScript framework, used in conjunction with back-end APIs to create dynamic and responsive user interfaces. Its inclusion supports more advanced front-end development needs, especially for intelligent building management platforms that require real-time data visualization and user interaction.

For the communication between components, FlaskAPI and Paho-MQTT are used. FlaskAPI is a lightweight Python web framework, which is used to expose services as REST APIs. This allows external applications or front-end interfaces to interact with the exposed service throught standard HTTP requests, making the system modular and easily extensible. Paho-MQTT facilitates real-time communication between components using the MQTT protocol. This enables data-intensive services to be processed by multiple consumers without overhead.

3.1.4 Data Acquisition and Datasets

For this project, no pre-existing dataset was utilized in the research process. Instead, the system is designed to autonomously collect and leverage real-time IoT data. Data sources include temperature and humidity sensors, device consumption metrics, and the overall energy consumption and generation profile of the building. The system operates by dynamically acquiring and processing information from these diverse IoT sources, enabling a comprehensive understanding of the building's environmental conditions, energy consumption patterns, and renewable energy generation. Data is collected and aggregated every 5 seconds and spans approximately 1 year prior. This approach ensures that the analysis and findings of this dissertation are based on up-to-date and relevant data, reflective of the current operational state of the monitored environment. Table 4 specifies the overall structure of the data produced in the community by the buildings.

Table 4 Data structure from acquired data

Type	Data Type	Devices
power	Double	Refrigerator, Water heater, Microwave, Dishwasher, Kettle, HVAC, Sockets and Lamps
voltage	Double	Refrigerator, Water heater, Microwave, Dishwasher, Kettle, HVAC and Sockets
current	Double	Refrigerator, Water heater, Microwave, Dishwasher, Kettle, HVAC and Lamps
doorOpen	Boolean	Refrigerator
state	Double	Refrigerator and Water Heater
humidity	Double	Refrigerator
temperature	Double	Refrigerator and Water Heater

3.2 Technological and Social Challenges

This section describes the technological and social challenges of this thesis. It addresses topics such as privacy risks associated with the collection and use of personal data, data security, as well as ethical considerations related to the application of AI models.

3.2.1 Data Privacy and Protection

In the landscape of intelligent communities, where interconnected technologies and intelligent buildings form the backbone of urban development, the need for robust encryption mechanisms is indisputable. However, the real challenge lies in striking a delicate balance between fortifying security and ensuring seamless usability (Kitchin, 2014). As residents and businesses seamlessly integrate into the intelligent communities, the challenge is to provide a secure environment without compromising user experience.

The proliferation of IoT devices, from smart thermostats to connected appliances, introduces a greater attack surface for cybercriminals. These devices often lack robust security measures, making them susceptible to hacking and data theft. Cyber threats have become increasingly sophisticated, hackers devise methods to infiltrate networks, exploiting vulnerabilities and gaining unauthorized access to sensitive data (Neshenko et al., 2019).

Privacy violations erode trust in the technology that underpins intelligent communities and intelligent buildings. Gaining meaningful and informed consent from individuals is essential for data processing activities. Without explicit consent, individuals risk having their personal information collected, analyzed, and used without their knowledge or authorization (Pathmabandu et al., 2023). Data protection regulations, such as (*General Data Protection Regulation (GDPR)*, n.d.) in Europe, play a critical role in safeguarding personal information, establishing guidelines for data collection, storage, and usage, ensuring that individuals have control over their data. According to the GDPR, individuals have the right to control their data and not have their transactions linked or tracked. Compliance with data protection regulations is mandatory for organizations that manage and process sensitive information.

Phishing, the act of deceiving individuals into divulging personal information through fraudulent emails or websites, remains a prevalent threat in the digital age (Zieni et al., 2023). In intelligent communities and intelligent buildings, phishing can represent serious problems, as in communities, there is a high number of individuals who may not be experienced in detecting these attacks and can be more inclined to click on links or open attachments from familiar sources, such as community notifications or building management systems.

Navigating the complexities of data security in intelligent communities and intelligent buildings requires a comprehensive approach that addresses both technical and human factors. Robust encryption mechanisms, layered security protocols, and regular security audits are essential technical safeguards. Simultaneously, fostering cybersecurity awareness among residents and building staff is crucial to prevent phishing attacks and ensure informed data sharing practices.

In this way, the proposed solution addresses the challenges of data privacy and protection through a dual strategy encompassing both infrastructure-level and service-level points. At the infrastructure level, the use of containerized environments, immutable operating systems, and network isolation policies ensure that each component of the system operates within securely bounded contexts, minimizing the risk of unauthorized access and facilitating system integrity. At the service level, the intervention of Caravels is much more reduced because in fundamentally depends on the actual development, however due to this characteristic the system does not only permits a tailored approach to each service but also it allows for regular updates for new solutions to be implemented.

3.2.2 Ethical and Social Issues

It is imperative that the collected data is not only secured but also used ethically and for legitimate purposes. The responsible handling of data ensures that individuals' privacy is respected, and information is utilized in ways that benefit the community as a whole. Striking this ethical balance requires not only robust technical safeguards but also a commitment to transparent and accountable data practices.

Personal differences in attitudes towards privacy and data sharing add a layer of complexity to the ethical use of information. In a diverse and interconnected community, residents may have varied perspectives on what constitutes acceptable data usage or even on what is sensitive data. Recognizing and respecting these preferences is essential for creating policies and frameworks that align with the values of the community (Y. Zhang, Qu, et al., 2023). This inclusivity can help to foster trust among residents and ensure that the community is open to the diverse needs and expectations of its population.

In (Emami-Naeini et al., n.d.), the survey participants manifest concern about potential discrimination. Diversity and non-discrimination should be core principles guiding the development and implementation of intelligent community technologies. Inclusivity in the design and deployment of smart systems helps prevent bias and ensures fair treatment for all residents. Whether it's in the development of algorithms or the deployment of services, a commitment to diversity and non-discrimination promotes an equitable community that serves the needs of every individual, irrespective of background or identity.

By empowering users with transparency and control over what is happing in their own systems, Caravels directly addresses the ethical and social challenges. Granting individuals visibility into data collection, processing, and usage not only supports informed consent but also reinforces user input, allowing them to make decisions aligned with their own values and privacy expectations. This helps to address different perspectives on data sensitivity, while mitigating concerns related to misuse, discrimination, or lack of accountability. Ultimately, by implementing ethical principles, such as transparency and modularity, into the infrastructure and the service design, the system fosters a community that respects individual autonomy and promotes equitable access and participation in the digital environment.

4 Implementation – Caravels

As humans, we live in societies, as a collective group, establishing connections, trading, doing favors, and relating to one another. This means that when we design a solution for an intelligent community and buildings, it should not operate in isolation; it must be connected with neighboring systems. For example, in a rural setting, if you run out of eggs or sugar, you might ask your neighbor to lend you some. Similarly, if a road is blocked, your neighbor might warn you about it. The proposed solution aims to replicate this kind of collaborative behavior, allowing systems to share resources and information with one another.

There are many solutions available for intelligent buildings, intelligent communities, or a combination of both. However, most of these solutions fail to account for individual users within a community. They tend to either consider the community as a whole or focus solely on a single type of user. This approach often results in a dehumanized and uniform solution that overlooks the uniqueness of each individual, making it difficult for community members to adopt these new technologies.

In Caravels, a community is built by connecting existing, independent building systems, each one being a caravel, as illustrated in Figure 5. In this way, Caravels is a System of Systems (SoS), both at the community level and at the building level, giving that IoT devices compose intelligent buildings, which in turn compose an intelligent community. This architecture grants several benefits both for independence and coordination across levels: (i) each IoT device/building operates independently but contributes data and actions to the building/community intelligence; (ii) each building system is individually managed but consider broader decision-making at the community level (iii) adding or remove services do not require entire architecture redesign; (iv) buildings are physically distributed, yet digitally interconnected.

Caravels draw its name from the Portuguese ships of the 15th and 16th centuries, which enabled maritime expansion through their robustness, adaptability, and capacity for both individual exploration and fleet coordination. Analogously, Caravels envisions a distributed computing environment for intelligent communities where individual buildings (caravels) maintain autonomy while contributing to a larger collaborative ecosystem.

Each caravel, encompasses all the hardware and software resources deployed within a building physical and digital space, including IoT devices, edge computing nodes, and cloud services. Individual buildings within a community each host an instance of the Caravels framework, responsible for managing local resources, interacting with users, and coordinating with other caravels.

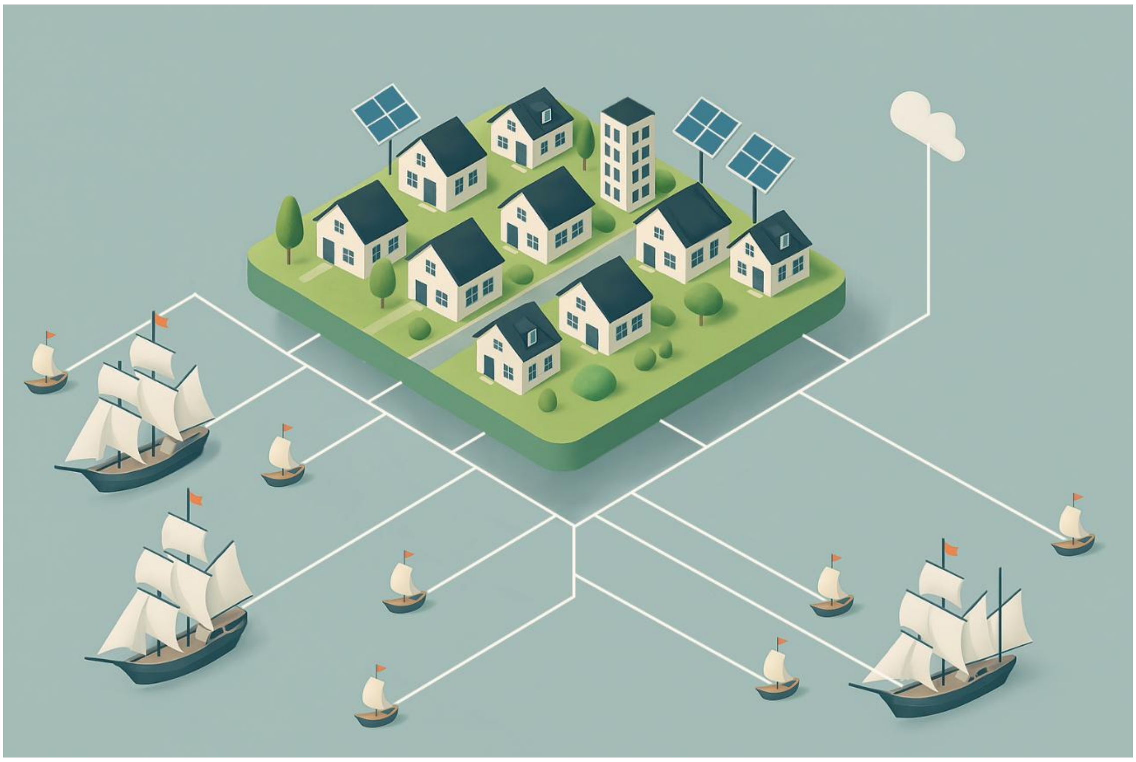

Figure 5 Caravels' interconnected buildings

The proposal outlined in this work aims to intelligently manage a building that is part of a community. This means that while the broader integration of the community is not ignored, the system implemented in each building is tailored to the specific needs of its users, creating a unique solution for those present in that building.

4.1 Architecture

Caravel's architecture was designed according to the following main requirements, regarding interoperability, scalability, privacy, and security:

- Support communication between each caravel (i.e., each entity/node part of the system, such as, the community operator, and a single building);
- Integration with heterogeneous IoT devices;
- Modular architecture to allow customization of services present in the caravel;

- Support for offline operation;
- Support for multiple host machines;
- Support host machine abstraction (e.g., single-board computers, laptops, and servers);
- Complete control over access to data is maintained by the data owner;
- The data owner has complete control over the stored data.

Caravels' architecture has a hybrid centralized-distributed model. Each node is able to perform computation tasks and play both passive and active roles, this means that this architecture supports both centralized and decentralized services. The community infrastructure has a central node, which handles all problems that affect the whole community, such as infrastructure operations or community-wide services. Every other communication or data transaction is carried directly between the nodes. Figure 6 represents the architecture of a Caravels node, however, it does not display arbitrary services.

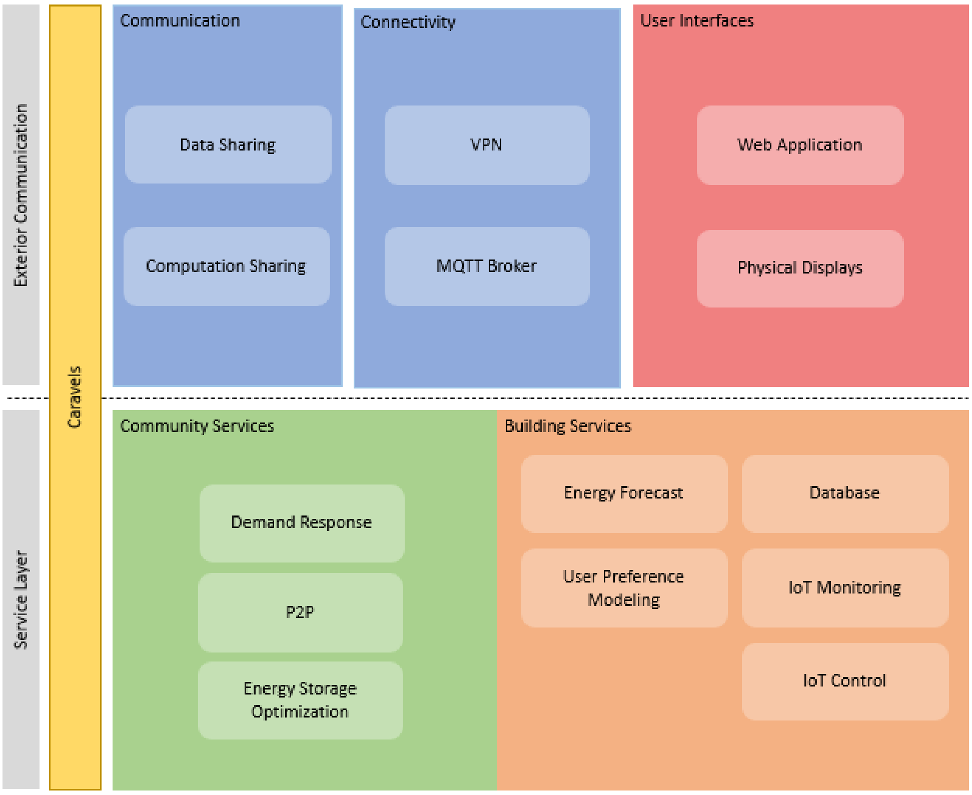

Figure 6 Caravels' node architecture

4.1.1 Containerized Infrastructure

A containerized infrastructure provides inherent advantages introduced by containers to the infrastructure such as the portability, scalability, or resource efficiency but also enables new benefits in the management of these containers. Containers encapsulate each functional component of the architecture, including communications, AI models, and visualization, into self-contained units that run consistently across different environments

One of the key advantages of containerization in Caravels is the ability to manage the lifecycle of individual components without affecting the rest of the system. Containers can be started, stopped, or replaced independently, supporting zero-downtime updates and fault isolation. Container orchestrators enable aggregated management across a cluster, not only across nodes but also across containers. This approach becomes increasingly significant when the number of containers deployed grows too much to manage them individually. For instance, Kubernetes provides monitoring and scheduling mechanisms that enable load balancing across containers. This ensures that the system avoids single points of failure, maximizes performance and resource efficiency. As failures can often disrupt the system, self-healing capabilities from the orchestrator automatically restart or redeploy failed containers, minimizing system downtime and simplifying management.

As applications may have traffic spikes or periods of low activity, it is important that the system can adapt accordingly by scaling vertically and horizontally. Vertical scaling refers to adjusting resources allocated to a single container, while horizontal scaling adjusts the number of containers responsible for dealing with a given workload. Orchestrators' scalability adjusts dynamically based on demand. Additionally, the flexibility of containerized supports multi-cloud or hybrid cloud deployments, allowing services to be hosted in the cloud or locally, depending on user requirements or constraints.

Due to container nature, applications run in isolated environments, and, at the cluster level, network policies can be implemented, the security in the infrastructure restricts access between system components. A container does not access the host machine hardware, and a node can only access containers explicitly defined by the system configuration.

To further enhance the resiliency of the system the host machines can adopt immutable operating systems. These are designed to run in a read-only state, this means that no changes can be made to the root filesystem and updates are done by replacing the entire system image. This greatly prevents unauthorized file modifications while increasing system consistency, facilitates rollbacks, and reduces configuration drift. These features collectively strengthen the security and robustness of the infrastructure.

4.1.2 Communication between Containers

Containers, by design, do not impose restrictions on how they communicate with one another. Communication mechanisms will depend entirely on the specific requirements of a service and the type of information being exchanged within a given use case.

For instance, when disseminating monitoring values across multiple services, the publish/subscribe protocol of MQTT is employed. It facilitates efficient one-to-many communication, allowing multiple services to simultaneously consume the same data stream. A practical example involves both the User Interface (UI) and the database service receiving real-time monitoring values. This approach avoids straining the data collecting and publication service, given that the system cannot scale that service or it would lead to duplication problems.

On the other hand, HTTP requests are utilized in processes that require a direct request/response or tasks involving active service participation. For example, triggering an action on the system on demand, which relies on HTTP synchronous communication capabilities.

By leveraging these complementary communication protocols, MQTT for asynchronous, event-driven data sharing and HTTP for request/response or action-oriented tasks, the system achieves both flexibility and efficiency. This design ensures scalability and reduces unnecessary resource consumption while maintaining support for different operational requirements.

4.2 Tools and Services

This chapter explores the tools and services used in this project. Such resource sharing mechanisms and task distribution across multiple systems, and an overview of developed services that support these operations.

4.2.1 Data and Service Sharing

Data sharing leverages the ability to share information from certain parts of the system with a given set of community members. For instance, this information can be the readings of a given IoT or the output of a determined service. A key technique to enable this feature is the use of tunnels.

A tunnel is a secure communication channel that transmits data or requests between endpoints, between machines, or over networks (Nasri et al., 2020). By encapsulating one protocol within another (e.g., secure shell or VPN tunnels), tunnels allow resources or services to be securely accessed and utilized across different machines or networks. In computation sharing, tunnels redirect workloads, making it possible for systems with limited resources to offload tasks to more capable machines.

In the context of a community, to be able to create a cluster where each member possesses a machine, they need to be able to communicate with each other. There are some protocols that can be implemented to connect machines to remote networks or machines to machines, such as VPN, Secure Shell (SSH) and application or container-level tunneling.

VPN-based Tunneling is used to provide a secure connection between networks, enabling seamless sharing of computation across geographically distributed systems, this technique allows for machines that are not in the same network to communicate with each other without having to expose themselves to the public internet. However, the same thing that presents itself as an advantage can also be viewed as a potential risk, because by connecting the whole machine to the given network, means that not intended access may take place, for example, if I want to share information directly with my neighbors with a VPN, their machine will have access to my whole home network.

SSH-based Tunneling is used to establish secure, encrypted connections between individual machines or systems, enabling remote access and data transfer over untrusted networks (Ylonen et al., n.d.). This method allows users to execute commands, manage resources, or forward traffic securely to a remote system without exposing sensitive services to the public internet. However, while SSH provides granular control over specific connections, its focused access can also introduce risks. Granting SSH access to a machine creates a pathway that, if compromised, could enable unintended lateral movement within a network. For example, if a user shares SSH credentials with a neighbor to grant temporary access to a specific machine, that neighbor could potentially exploit misconfigurations (e.g., weak firewall rules, unrestricted port forwarding) to access other devices on the same local network.

Both tools described above involve network-level access, which connects entire networks or individual machines, where the risk, in the context of this project, is overexposure. Mitigating this problem implies that the attack surface needs to be reduced. By exposing only specific services, in the form of containers and applications, external access will be restricted to predefined ports, protocols, or endpoints, decoupling the determined exposed service from the hosting infrastructure. To this intent, tools like Ngrok[14] and Fast Reverse Proxy (FRP) as well as role-based access control features provided by Kubernetes favor minimal exposure and align with zero-trust principles, where no implicit trust is granted to users or systems.

By employing tunneling mechanisms, the system achieves optimized resource utilization through task redirection to capable nodes, scalable architecture via dynamic integration of remote resources, and enhanced collaboration among distributed participants. However, the broader network access inherent to tunneling imposes the necessity of safeguarding each network. Balancing the approach between network and application-level tunneling ensures efficient use of resources and threat minimization. Figure 7 provides an illustration of the data or service sharing communications between nodes in Caravels. Data produced by Service A, hosted on Node 1, is consumed by Service B, hosted on Node 2.

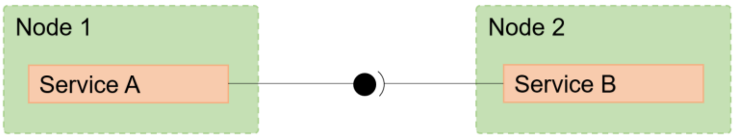

Figure 7 Data or service sharing communication in a Caravels

4.2.2 Physical Mobility

This section describes how physical mobility, and dynamic orchestration of workloads across geographically distributed or community-shared resources, enable decentralized computation sharing.

[14] https://ngrok.com/

Computation sharing leverages the ability to distribute computational tasks across multiple systems. This technique also uses tunnels to provide secure connections between systems. Figure 8 illustrates Service A being deployed, by Node 1, on the Node 2, and avoiding communication with other Node 2 services while preserving communication with Node 1 services.

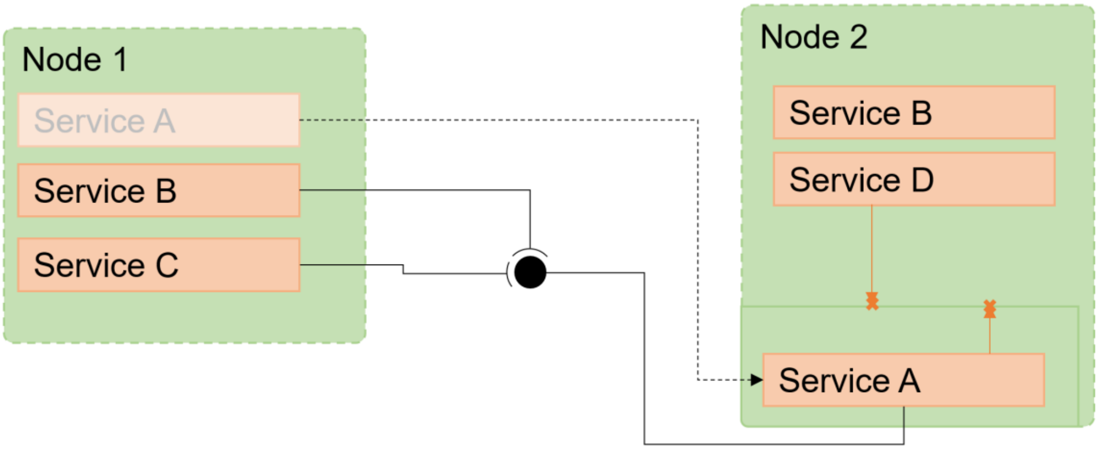

Figure 8 Computation sharing diagram

Physical mobility, in this context, refers to the seamless relocation of computation tasks in different nodes, in a community, dealing with resource availability or cost constraints. By utilizing Kubernetes capabilities of treating dispersed resources as a unified computation resource, such as persistent volume migration and automatic pod rescheduling, the system enables workload portability.

Community members contribute with any computation-able machines (e.g., edge servers, personal computers), which they can share to the community computation pool or not. Doing so, Kubernetes schedules tasks to these nodes based on real-time resource capacity, specialized hardware, or energy efficiency.

By optimizing resources owned by the members of a community, the community reduces reliance on centralized cloud providers and members the dependency on cloud services, as well as regaining data sovereignty, which opens doors to data sensitive applications which gain end-users confidence, because the data is self-owned, and if cloud-based would face compliance with regional data protection laws, which could impact the feasibility of determined features.

4.2.3 Service Deployment Process

To deploy a service, several steps must be followed. As illustrated in Figure 9, these steps encompass service configuration from the development phase through to the deployment phase. Additionally, considerations must be made regarding the location of the service store prior to deployment, ensuring it is accessible to each building's system.

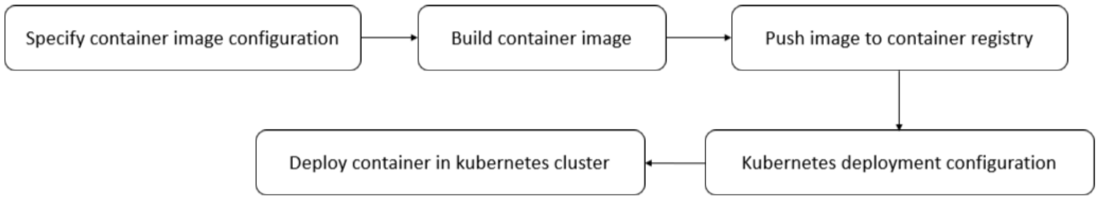

Figure 9 Service deployment process diagram

Deployment of a service starts with specifying the container image configuration, which defines the template that deployed containers will replicate. Containers have internal and external ports. By specifying in what port the encapsulated software communicates, configuration allows the host machine to route traffic coming from the external exposed port to the container's internal port. For example, a host machine running two services, both communicating on port 80 internally, can distinguish traffic by specifying different external ports to each container, ensuring proper routing. In addition, configuration also defined dependencies, including the software's programming language, required packages their respective versions. These dependencies guarantee consistency and compatibility across different deployment environments.

At the image building stage, the target architecture must also be specified. For example, a Raspberry Pi 4 typically uses arm64 architecture which will be different from conventional laptop, which has amd64 architecture. This foresight regarding deployment hosts ensures that the container images are compatible with the hardware they will run on. Once images are built, they are pushed to a container registry, which can be local or cloud-based. Figure 10 shows a printscreen of a developed web application responsible for building an image for a given service.

As mentioned earlier, all containers in this project are deployed within a Kubernetes cluster, requiring more configuration. Kubernetes offers a wide range of options, and while configurations may vary between applications (e.g., some services might use default settings), this thesis will only cover the ones relevant to the project.

In a Kubernetes cluster, there are three configuration objects essential to deployment: Deployments, Services, and ConfigMaps.

A Deployment object defines the number of replicas the system for scaling under heavy traffic, what image will be deployed, and which nodes can run the specified container (e.g., specifying that service A should only run on node A). If such constraints are not defined, Kubernetes automatically optimizes container allocation across the cluster.

A Service object handles network communication for containers. Rather than connecting directly to individual containers, traffic is routed through the Service object, which ensures that requests are directed to a functioning container. This abstraction simplifies communication and improves fault tolerance.

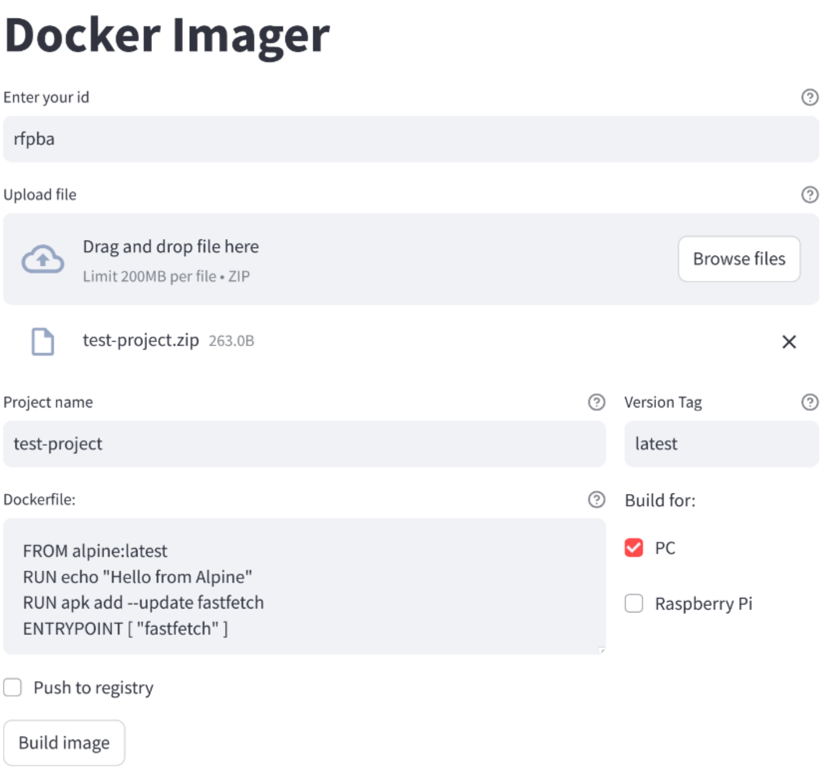

Figure 10 Printscreen of a web application for building container images

ConfigMaps store configuration data. To understand how they impact the system it is necessary to understand how a Kubernetes cluster works, containers can be deployed in different nodes and the system can have the same image for different purposes, so this leads to the problem of how to configure this deployment across the cluster. Instead of having a configuration file saved on each machine, which becomes difficult to manage (believe me, I have done it), configuration is stored in the ConfigMap object making it accessible anywhere in the cluster.

Establishing a well-defined deployment process enables the creation of a Continuous Integration/Continuous Deployment (CI/CD) pipeline, which automates and simplifies the deployment workflow. A CI/CD pipeline ensures that updates to the codebase are seamlessly built, tested, and deployed, reducing manual intervention and minimizing the risk of errors. Enhancing reliability, accelerating deployment, and supporting iterative development practices.

4.2.4 Available Services

Following the system architecture, a set of specialized services has been developed to serve different domains, both at the individual building level and for the broader community. These services are structured into four distinct categories:

- Infrastructure – This foundational layer handles essential system operations, including communication protocols, data exchange mechanisms, and container management;

- User Interface – Responsible for providing intuitive graphical tools, this layer enables interaction between users and the system, ensuring accessibility and ease of use;

- Domain – As the core of the architecture, this layer delivers key functionalities necessary to achieve the system's objectives. It integrates various domain-specific services to support the proposed solution;

- Artificial Intelligence – Closely related to the feature layer, this component leverages advanced models to enable data-driven insights. It handles model lifecycle steps after deployment, which include data processing, model training, evaluation, and monitoring.

At the infrastructure level, the developed services are closely integrated with either the container orchestration tool or physical components of the system:

- Computation Sharing: Enables members to allocate a determined container on another member's machine. The output from the given container can only be accessed by the service owner, not the owner of the machine;
- Data Sharing: Members can share information with another member, a defined set of members, or, ultimately, with the whole community. It can be the output of a service or data generated by an IoT;
- Service sharing: Similar to data sharing but allows for bilateral communication, meaning that other members can not only receive data but also interact with the shared service;
- Service Adaptability: Manages the services running on a member machine, enabling them to add or remove services as needed;
- IoT Control and Monitoring: Monitors and controls IoT devices owned by the member. It also makes collected data accessible to any service on the node in real-time;
- Storing: Responsible to manage database systems, ensuring redundancy, and well-functioning.

For the user interface category, three services were developed:

- Community Interface: Is a web application that serves as a way for the community to monitor the community energy components, trigger actions, and analyze outputs provided by the community services;
- Building Interface: Web application that provides information and enables user interaction with the services. Information displayed will depend entirely on services available at the node;
- E-Paper Display Interface: Shows energy monitoring data, such as energy consumption, and generation. It was specifically developed to work with a 2.7 inch e-paper display compatible with a Raspberry Pi 4.

There is an additional distinction between domain services, community-wide services and building-level services. Community-wide services imply that there is some type of communication and organization between members and affect the whole community, whereas individual services outcome only affects a given member and does not depend on other members.

Developed community-wide services are composed of energy-focused services namely:

- Energy Storage Optimization: Manages the energy storage of the community, based on linear optimization. Schedules batteries charges or discharges, based on energy price, energy demand and generation;
- Demand Response: Responsible to identify opportunities for events, classify eligible participants sent by community members, and manage and monitor events;
- Peer-to-peer Bidding: Based on auction model, calculates all possible transactions between members. Depends on energy surplus and buying and/or selling prices defined by each member.

Developed building-level services, essentially, act on user comfort, and energy efficiency:

- HVAC optimization: Automates HVAC systems in a defined room, based on past user interactions. Depending on temperature, humidity, and user location of the room;
- Energy Forecast: Predicts consumption, generation and flexibility of each IoT device for the next day;
- Indoor Location: Tracks a given user location inside the building, to provide additional environmental context;
- User Preference: Manages user preferences under different contexts, to enable more intelligent automations.

The Machine Learning category houses all the models utilized by a given node:

- Energy Consumption Forecast: Predicts energy consumption of a building for the next day, in periods of 24 hours corresponding to each hour of the day;
- Energy Generation Forecast: Predicts energy generation of a building for the next day, in periods of 24 hours corresponding to each hour of the day;
- Preference Learning: Applies a reinforcement learning model to learn a given value for a given preference based on interactions with the system and user feedback.

4.3 User Preference Tree

This section focuses on the tree-based modeling used to create and store user preferences, which evolves over time with user interactions through the inclusion and addition of conditions and contexts to the structure.

4.3.1 Tree Structure and Components

The proposed structure is based on a hierarchical tree structure, with the user as the root node. This structure organizes user preferences into distinct layers, enabling modularity, adaptability, and scalability. To this end, as seen in the example tree structure in Figure 11, five node types are available for user preference modeling:

- User node: represents the root of the whole structure of a given user. Acts as the aggregator for multiple preference nodes;
- Preference node: defines a specific category of preferences, such as temperature, lighting, or sound settings. It serves as a preference root for values and any related conditional logic. Each preference node can branch into one of these nodes: value or condition;
- Condition node: introduces a logical rule that determines appropriate context or value under circumstances defined in the condition function. For instance, a condition could account for environmental variables, such as time of day or season. Every condition node is required to have a default output, ensuring robustness in cases where conditions to other context nodes are unmet. A single condition node may branch into multiple context nodes, each representing a different scenario;
- Context node: these nodes are always preceded by condition nodes and define outcomes based on the parent condition. A context node may either contain additional condition nodes, facilitating multi-level logic to construct complex scenarios or provide more granular control, or directly store a value node that represents the output value for the corresponding context;
- Value node: stores the actual value of a preference. These values can take various forms (e.g., single values, intervals, or lists) and even evolve from one form to another. In advanced implementations, a value node may incorporate machine learning mechanisms, such as reinforcement learning to dynamically adjust based on user interactions.

Through the use of condition and contexts nodes, the structure can store conditional logic. It enables the system to dynamically adapt user preferences to the user context, for known variables that can be retrieved or calculated. By implementing context ingestion, the system collects and integrates external variables into its decision-making process. These variables could include environmental data, user behavior patterns, or external sensors. The condition nodes interpret this context and map it to corresponding context nodes, which will define on the corresponding leaf node the appropriate value node.

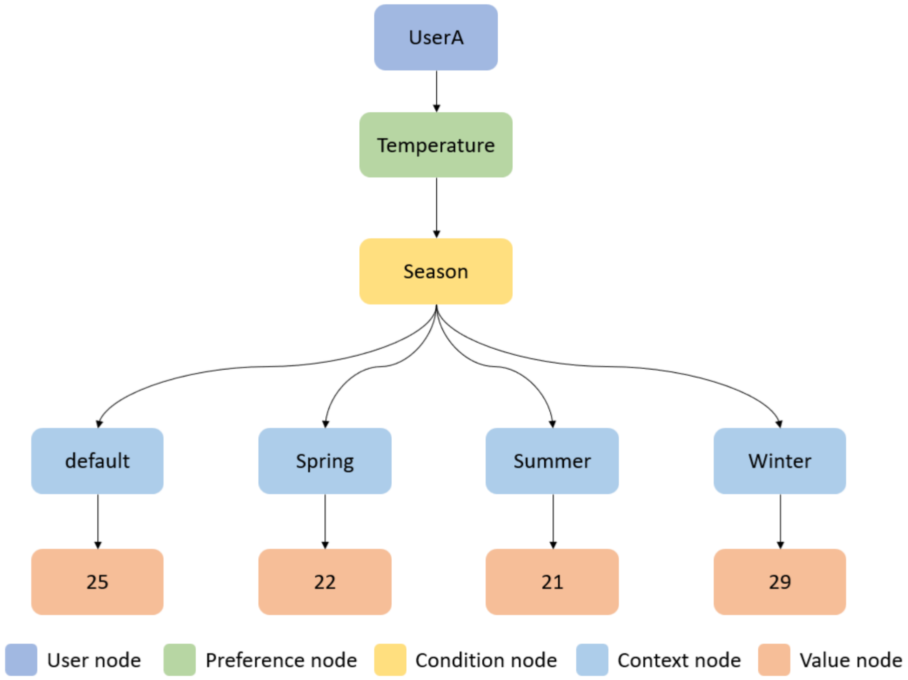

Figure 11 Example Tree Structure

4.3.2 System Interaction

Initially, user structure trees only contain the user node, so for the structure to grow and have a meaningful purpose preferences nodes need to be added. From the catalog, by adding nodes to their structures, as can be observed in Figure 12, the structure will fit their requirements or needs. Once the structure hits a certain level of complexity, it will increase its level of personalization and uniqueness to that given user.

However, as the system seeks integration with different services and knowledge exchange between users, problems arise. For a service to use a preference from a structure, it needs to know what preference it should refer to, for example, a service that controls the temperature of a room needs to know the "Temperature" preference of the user. If the user has multiple preferences, what is the preference it is going to pick? To this end, the system defines a catalog of preferences from which the user can add, and the service can select what is the preference it is going to work with, and which will search in the preference nodes present in the user structure.

A similar approach is also applied to the condition nodes, although for a different reason. To enable knowledge exchange between users, a correlation between the two structures needs to be set. By providing a catalog of conditions, the system ensures consistent behavior of a condition across structures, avoiding conflicts between conditions with the same objective but producing different contexts.

At a stage where there are multiple preference values, the user might want to define a new value for this node or even change its type. To enable this action, value nodes are able to mutate dynamically, preserving their relations and positions in the structure.

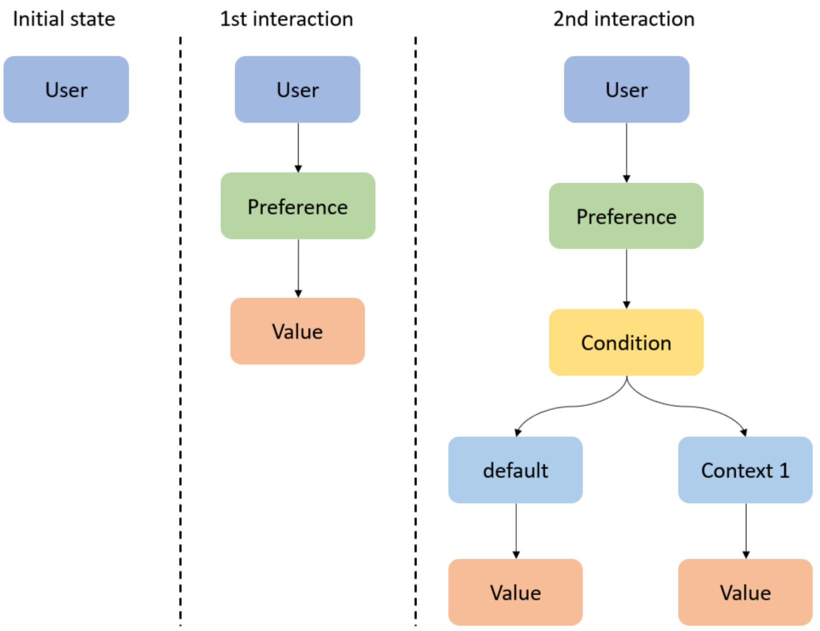

Figure 12 Structure evolution through interactions

4.3.3 Organic Growth

The user preference tree is designed to evolve, growing organically either by user individual interactions with the system or by having contact with other users. This dynamic evolution ensures that the user preferences stored remain aligned with the user's wants and needs while fostering collective intelligence through decentralized learning. Apart from the owner of the structure, it also needs to interact with services from the building and with structures of other users. To support different modifications and interactions with the structure, a set of operations were defined, as shown in Table 5.

Table 5 Possible actions with the user structure

Action	Description
Planting	Addition of a new node to the structure
Weeding	Deletion of a given node
Harvesting	Extracts the current value of a specified preference
Trimming	Cuts a part of the overall structure
Grafting	Merges two structures
Flourishing	Interacts with the structure, directly or indirectly, to modify the value nodes

It is possible to identify two mechanisms of growth, Human Action-Reaction (HAR) adaptation or Human-to-Human (H2H) adaptation:

- HAR adaptation: defines two types of feedback, implicit and explicit. Implicit feedback refers to the definition of user preferences through behavioral patterns, and explicit feedback to the user manually modifying the tree structure;
- H2H adaptation: preference propagation enables users to share and adopt subsets of their preference tree with each other.

The organic growth of user preference tree is sustained by technical frameworks that capture individual and collective behaviors, adaptively refining decision logic while preserving privacy and scalability. These refer to interaction tracking, versioned tree evolution, and privacy-preserving growth.

In interaction tracking, organic evolution begins with granular event logging, which records user decisions, such as manual overrides, and task delegation choices. These logs serve as the raw input for detecting behavioral patterns which can then be used by machine learning models, such as clustering and collaborative filtering, to analyze interactions to propose structural modifications to the tree. By correlating frequent overrides with the context of the current moment, the system infers unmet user needs and adjusts conditions or thresholds autonomously.

To support incremental refinement, the tree undergoes version-controlled updates. Historical snapshots are archived, enabling rollbacks to prior states if new branches introduce conflicts or inefficiencies. For instance, if a user's preference for "energy-efficient nodes" gains sub-branches based on time-of-day patterns, only these new nodes are transmitted—not the entire tree. This lightweight approach allows seamless integration of community-shared preferences without disrupting ongoing operations.

Structures are stored and owned by the user. Meaning that their structure is not spread within the community, without explicit intent of the user. This allows users to share and evolve their trees, forming collective wisdom, without exposing sensitive information. This topic will be further developed in the next section.

By combining the functionalities provided by structure interaction with privacy-aware collaboration, the preference tree remains both personally relevant and community-aligned. Versioning ensures stability amid experimentation, while collective knowledge prevents stagnation by integrating dispersed insights. Over time, the tree transitions from a static rule set to a living artifact, continuously shaped by its users' evolving priorities and shared knowledge.

4.3.4 Structures Merging

This section addresses the mechanism by which individual preference trees dynamically merge into contextual collective structures when users co-exist in shared environments. The merged structure operates as a transient, consensus-driven hierarchy that reconciles individual preferences while preserving autonomy and enabling collaborative decision-making without altering users' individual preference trees.

The action of merging structures refers to the creation of a temporary preference tree derived from the hierarchical rules of multiple users interacting in a shared space. This mechanism keeps individual trees intact for personal contexts, and merged trees only govern shared environments, dissolving when users exit the context.

The purpose of structure merging is to enable seamless collaboration without requiring permanent compromise of personal preferences and balance competing priorities, but because user structures are unique, they can have different organization, nodes and values. To merge structures, they are represented as directed graphs, with merging algorithms identifying isomorphic subgraphs and divergent paths. For these divergent paths, conflict resolution strategies are applied, which consider:

- Weighted Strategy: Users assign priority scores to the path, with aggregated weights setting the dominant path;
- Contextual Arbitration Strategy: Users add a preceding condition node, defining the primary path based of the current context defined by the condition;
- Majority Voting Strategy: Users vote to select the dominant path.

This framework ensures merged structures act as adaptive, context-sensitive intermediaries, harmonizing individual and collective needs without compromising personal preference integrity.

5 Case Studies

In this chapter, the proposed solution and its features are presented. To test and validate, several case studies were developed. Table X shows their description and objective.

Table 6 Proposed case studies

Name	Main Objective
Caravels technical testing and validation	Evaluate Caravels' technical features to integrate services with communities
Users Virtualization	Assess the system's ability to understand and adapt to user preferences, using a graph-based structure
Multi-User Implementation	Explore the modeling of preferences in shared contexts, considering conflict resolution strategies, group decision-making frameworks, and context-aware negotiation mechanisms
Intelligent Energy Community powered by Caravels	Implementation of Caravels in an Energy Community, integrating several energy models

5.1 Caravels Technical Testing and Validation

This case study aims to evaluate Caravels' technical features to integrate services with communities. It consists of three parts: (i) service adaptability, where the system calls a new service and, seamlessly, changes services with the same purpose; (ii) computation sharing, deployment of service in a system different from the one that uses its outputs, and (iii) information sharing, where a system shares information from an IoT device or a service to other community members. This case study is based on the case study presented in one publication:

- [Conference] **Rafael Silva**, Luis Gomes, Zita Vale (2024) "Caravels: a Decentralized Container-Based Infrastructure for Sustainable Human-Centric Intelligent Energy Communities", presented in 2024 International Conference on Soft Computing Models in Industrial and Environmental Applications (SOCO 2024), doi: 10.1007/978-3-031-75013-7_23.

In the case study's first part, only one user will be considered as the feature only affects the user's system. To enable service adaptability, the service in question is considered reachable by the members' system. This means that the service can either be in a public or private registry. This enables the user to pull services that are publicly available or in registries owned by the community or by the user itself.

A system may have different services at different times, meaning that it is not static, similarly to requirements that are dynamic over time. Figure 13 illustrates the services available at "Node α" over three different moments. Initially at T0, the node in question only possesses two services. The Storage service is responsible for storing the data coming the Consumption Monitoring service, by subscribing to the MQTT topic in which it publishes. The Consumption Monitoring service communicates with each IoT device capable of reading energy consumption, it provides both individual and aggregated consumption of the selected group.

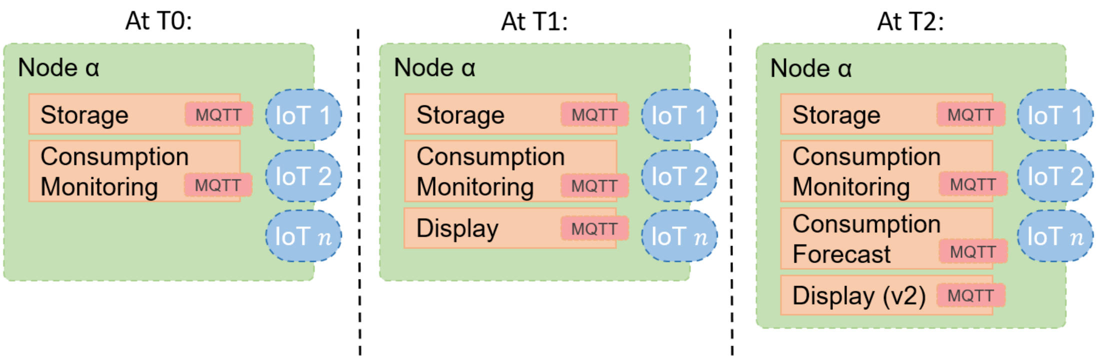

Figure 13 Services present in Node at three moments (T0, T1, and T2)

For this case study, a display service was developed. This display service is hardware dependent and is only compatible with a Raspberry Pi and a specific e-Paper display, because it communicates directly with predefined pins. By following the process illustrated in Figure 14 and assuming that the user possesses a compatible e-paper display, the user can call a new service to be deployed on its system.

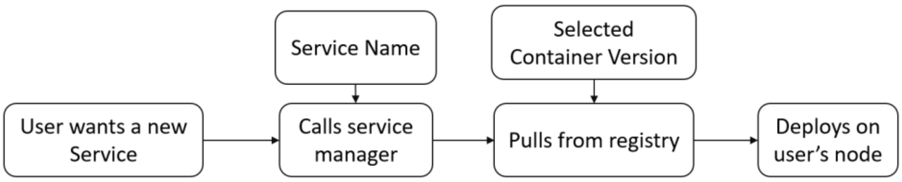

Figure 14 New service integration process

Following this deployment, the display service follows a similar approach to data input. Both consume the same MQTT topic. This ensures that there is less data traffic, avoiding straining the Consumption Monitoring service in this case, and that multiple services can consume the output data without modifying the Consumption Monitoring service. As seen in Figure 15, at T1, prior to the deployment of "Consumption Forecast" service, only the current consumption is displayed. At T2, after the deployment of the service, both current and forecast consumption are presented.

a) b)

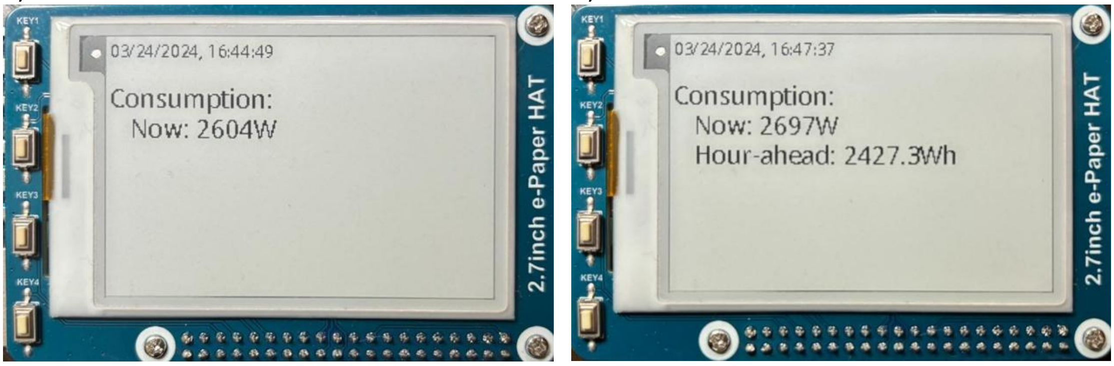

Figure 15 e-Paper display: a) at T1, b) at T2

The second part of this case study will delve into the concepts of data and service sharing within caravels. Each IoT device is continuously monitored, and its data is streamed through a specialized MQTT broker service, which acts as a central hub for internal communication. This configuration allows a variety of applications and services to access and consume the data generated by these devices.

In exploring these concepts, we can identify two separate key approaches: data sharing and service sharing. Data sharing is primarily concerned with the dissemination of output data, allowing systems to receive information without the ability to interact. In contrast, service sharing enables a more dynamic form of communication, where two systems can engage in bilateral interactions. This means that not only can data be retrieved from the shared service, but other systems can also send commands or requests, thus facilitating a more integrated and responsive network of services and devices.

To provide a comprehensive demonstration of both features, as depicted in Figure 16, three distinct services were deployed within a two-node system. In Node A, a data broker and a monitoring service were established, facilitating direct communication with a weather station installed in the respective building. This weather station collects real-time meteorological data, such as temperature, humidity, and atmospheric pressure. The data broker service aggregates the data corresponding to the weather station from the monitoring service and streams it to the subscribed services. Node B, on the other hand, has only a weather forecast service deployed. This service uses the real-time data received from Node A's weather station to generate short-term and long-term weather predictions.

The interaction between the two nodes is reciprocal: Node A shares the raw data collected from its weather station so that Node B can provide accurate forecasts. In return, Node B disseminates the forecasted data back to Node A, allowing it to offer enhanced insights and analytics based on both real-time conditions and predictive models. The data and service sharing process enhances the overall functionality of a community by combining resources of each member, whether they are hardware or software.

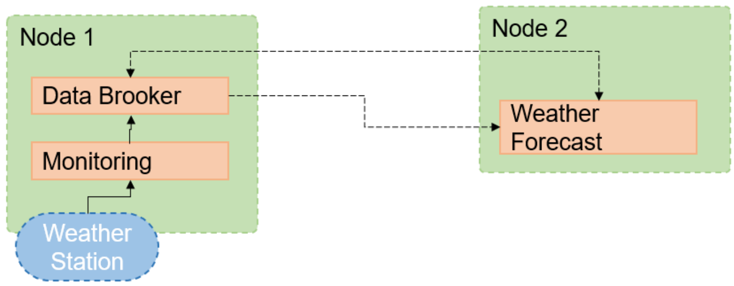

Figure 16 Data and Service Sharing case study

The third part of this case study combines the features already mentioned, by utilizing the dynamic integration of services in a node and the data exchange between nodes. Computation sharing regards the exchange of computation power of one member to another. This means that a member's system will request another member's system to host a given service. Because this computation unit is not to be moved physically, the hosted service will exclusively communicate with the source node services, meaning that it will be executed in complete isolation, avoiding communications with the other services of the node.

Developing on the scenario presented before, as Figure 17 shows, Node 1 has installed photovoltaic (PV) panels, but its system runs on a Raspberry Pi 4 which does not have enough computation power and resources to run the Solar Generation Forecast service, because it requires a Graphical Processing Unit (GPU). As Node 2 runs on a server equipped with a Nvidia GPU and has enough computation power, the service will be hosted by Node 2 system in isolation. Node 1 aggregates the data produced by the "Weather Forecast" service with the Solar Generation Historic and injects it into the "Solar Generation Forecast" service for inference tasks.

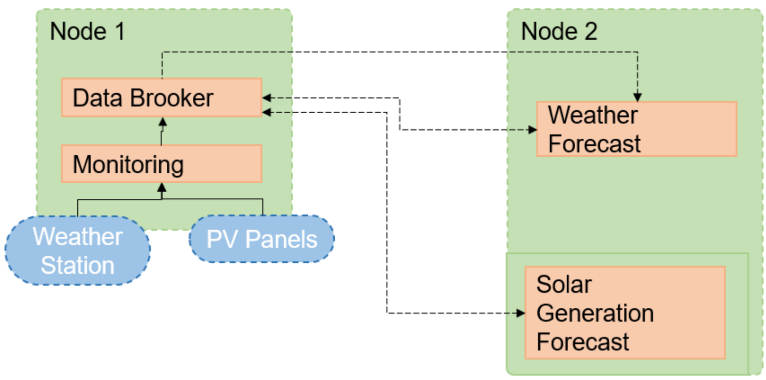

Figure 17 Computation Sharing case study

The results of this computation sharing are shown on the web application, present in Figure 18, corresponding to the interface of Node 1.

To conclude, based on the different requirements of each household, the system grants a certain degree of adaptability only limited by existing developed services. Containers are independent of each other, except for communication, so the node only needs to deploy a container responsible for a service, removing direct dependencies between containers, and allowing for different implementations for the same service. The system can provide some degree of portability given that in the event of failure, to get the system running again only the assignment of services on the new node host is required. Additionally, through data and service sharing, as well as computation sharing, Caravels is able to mitigate both physical and computational resource constraints.

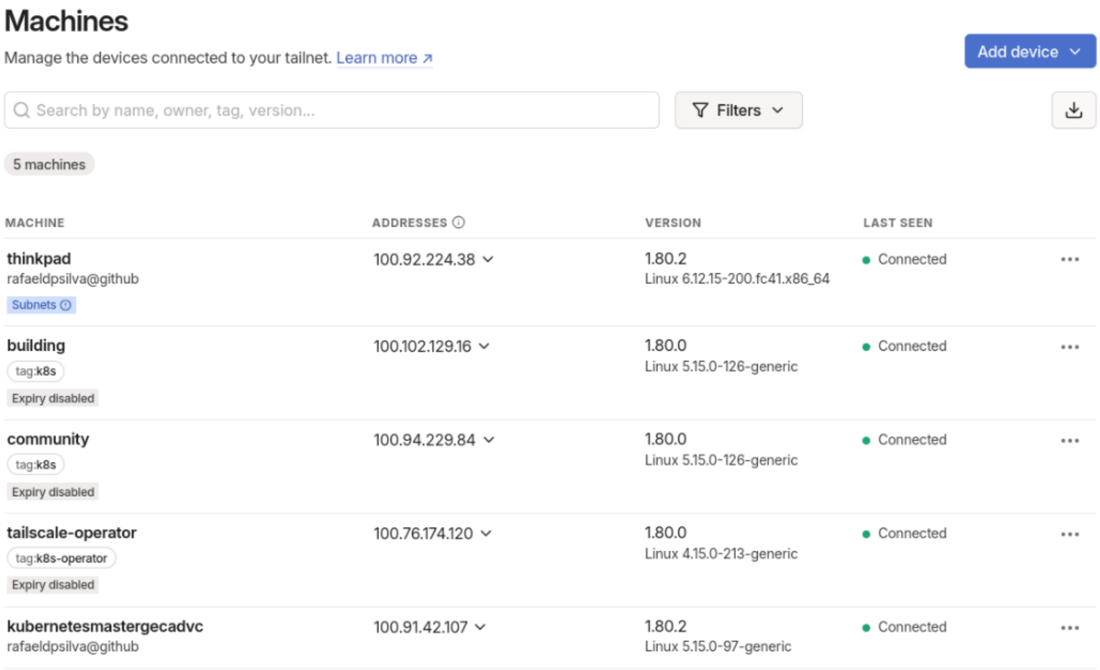

Figure 18 Tailscale Interface (print screen)

5.2 User Virtualization

This use case aims to assess the system's ability to understand and adapt to user preferences. The goal is to create an interface that benefits both the user and the system. It allows the user to record their preferences in a dynamic structure that evolves over time. For the system, it allows services to interact with the user preferences and to have these preferences adjusted to different contexts and environments while maintaining abstraction and without needing direct user feedback. This case study is based on the case study presented in one publication:

- [Book Chapter] Rafael Silva, Bruno Ribeiro, Luis Gomes, Zita Vale "Home Energy Management Models", accepted in "Home Digital Twins".

To facilitate interaction with the user, a user-friendly visualization graphical was created, as seen in Figure 19. It is able to perform several actions on the user structure, such as:

- Planting: Adds a node to the user structure;
- Weeding: Removes a node from the structure;

- Modifies the condition nodes existing on the preference sub-structure, as well as their order. Adding or removing conditions changes the context to which the structure reacts;
- Flourishing: Only changes value nodes, modifying their value or type.

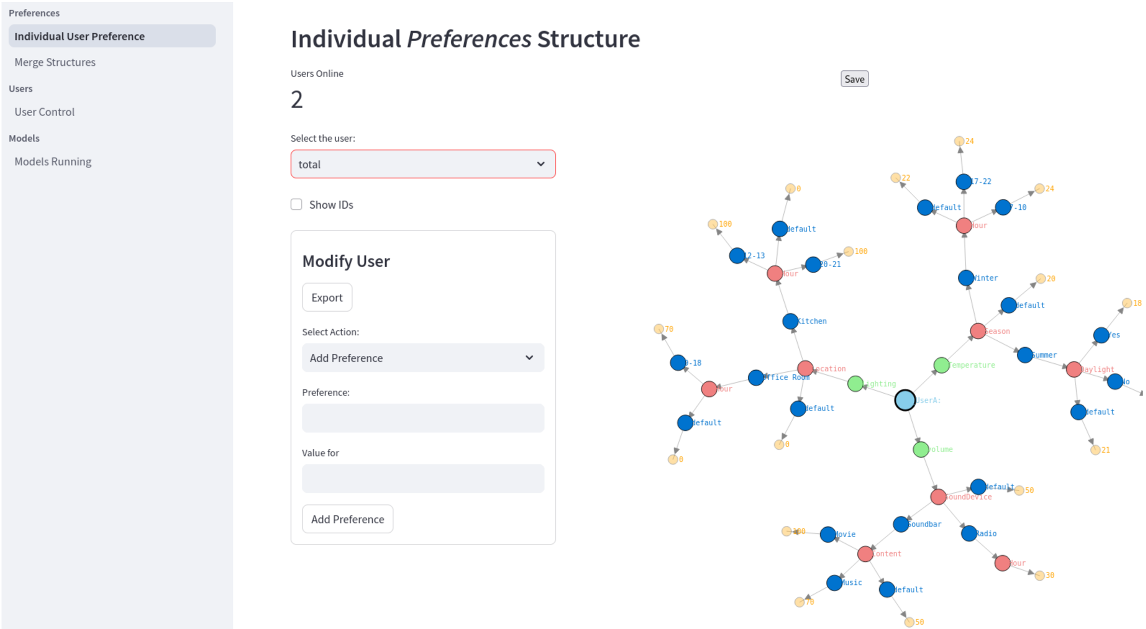

Figure 19 User Interface to interact with preferences (print screen)

At the initial stage of the user structure creation, the only node present in the structure will be the User node, without any virtual functionality. To enhance this structure, three preferences were added:

- Temperature: Enables the system to perceive and respond to ambient temperature variations, allowing for automated adjustments based on predefined conditions;
- Volume: Facilitates the detection and regulation of sound levels, adapting output to optimize user experience and environmental conditions;
- Lighting: Adjusts illumination intensity dynamically in response to ambient light levels, ensuring optimal visibility and energy efficiency.

When associating a new preference to the user's structure, each preference must have a corresponding value node. Modifying the value node type introduces some flexibility. For example, in the temperature preference, the type is set to Interval, ranging from 23ºC to 25ºC, allowing for minor fluctuations in room temperature. Also, in Lighting preference, the type is set as List, switching between 70% or 100%. As seen in Figure 20 a), this creates a straightforward structure.

However, preference values remain the same even when the context changes. To introduce contextual flexibility, condition nodes must be implemented into the structure. In Figure 20 b), a new condition node is implemented on the Temperature preference, making the current value depend on what the current season is.

a) b)

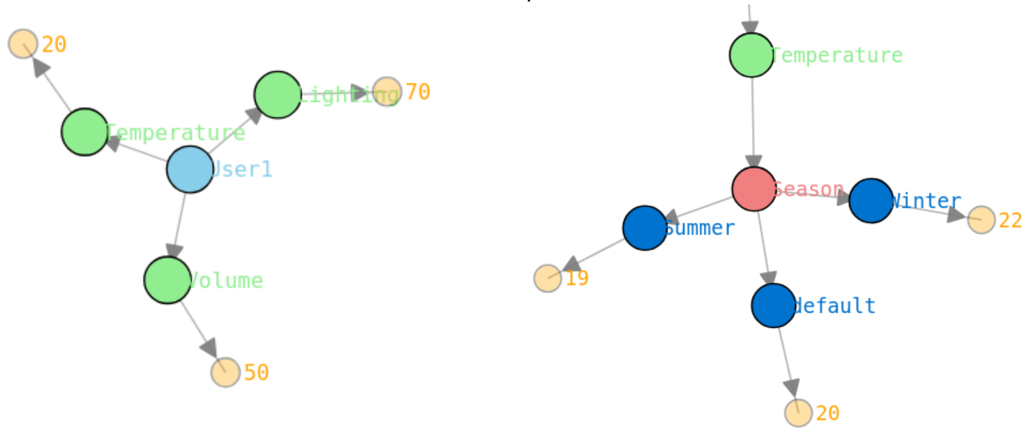

Figure 20 a) Whole user structure b) Evolution of Temperature preference sub-structure

Additional condition nodes are implemented using the same approach, which involves creating new context nodes and establishing a default fallback option. In some cases, the value node type is adjusted to accommodate more flexible values. These modifications are illustrated in Figure 21, presented in a graph for visualization. The goal of adding or modifying nodes within this user preference structure is to ensure that preferences adapt dynamically based on contextual factors for the following scenarios:

- Temperature: Seasonal variations play a key role in shaping user preferences. In warmer months, the user prefers cooler temperatures, with variations based on sunlight exposure, translating to lower temperatures in sunlit areas and higher temperatures in shaded ones. In winter, the temperature preferred is influenced more by the time of day, with notable differences between mornings, late afternoons, and evenings;
- Lighting: This preference is centered on artificial lighting, with an emphasis on energy efficiency. In the office room, the user prefers a comfortable lighting level, while in the kitchen, brightness should be sufficient for cooking tasks. If adequate natural light is available, artificial lighting is minimized to reduce energy consumption;
- Volume: The ideal volume setting depends on both the device’s speaker capacity and the listening environment. When audio serves as background noise, the user prefers a moderate volume. However, when media consumption is the primary focus, a more immersive sound experience is desired.

To further enhance adaptability, an RL Value type node can be introduced into the user structure. Unlike static or predefined values, this node enables dynamic learning and adjustment based on user behavior and contextual feedback. By leveraging reinforcement learning, the system continuously refines preference settings over time, optimizing user satisfaction while responding to changing environmental conditions.

This approach enhances personalization in scenarios where user preferences are not easily defined by fixed rules. For example, instead of relying solely on predefined temperature intervals or lighting levels, the RL Value node learns from interactions, adjusting preferences based on patterns in user behavior. Over time, this enables the system to make more accurate and context-aware adjustments, improving overall responsiveness and automation.

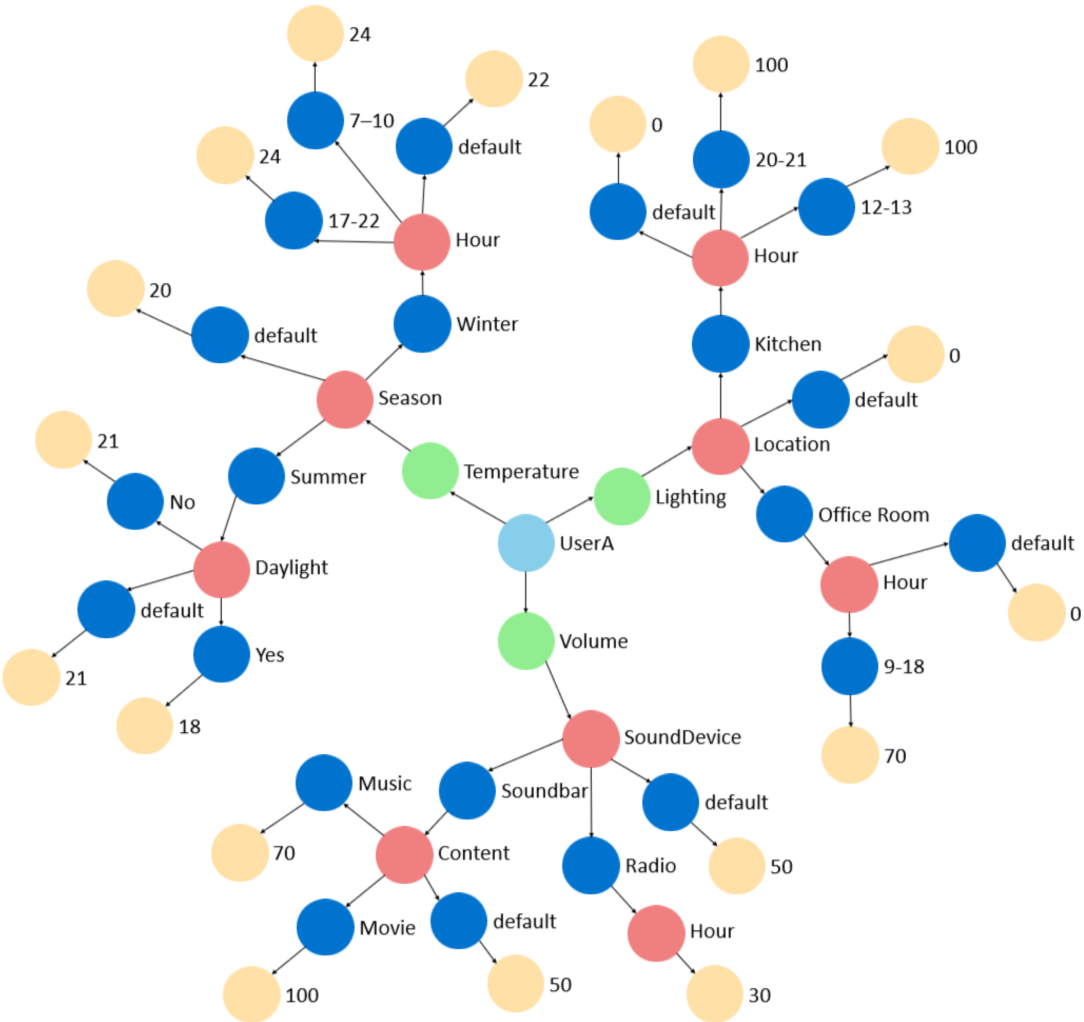

Figure 21 Whole user structure graph

Especially in user-facing systems such as audio playback, personalization can have several impactful factors. The RL model was implemented using a tabular Q-learning algorithm designed to learn on a custom environment that simulates multi-model contextual features such as device type, content being played, time of day, and user location.

The task of selecting an appropriate audio volume level is modeled as a MDP, characterized by the tuple (S, A, P, R, γ), where:

- S denotes the finite state space;
- A denotes the finite action space;
- $P(s' \mid s, a)$ is the state transition probability;
- $R(s, a)$ is the reward function;
- $\gamma \in [0,1]$is the discount factor.

The state $s \in S$ is a four-dimensional vector encoding the current contextual configuration:

- $s1$: Device type (e.g., phone, speaker) - 8 discrete bins;
- $s2$: Content type (e.g., music, podcast) - 8 discrete bins;
- $s3$: Time of day (e.g., morning, night) - 6 discrete bins;
- $s4$: Location (e.g., kitchen, bedroom) - 3 discrete bins.

Thus, the total state space cardinality is:

$$|S| = 8 \times 8 \times 6 \times 3 = 1152$$

The action space $\alpha \in A$ corresponds to discrete volume levels uniformly distributed over the interval [0.0, 1.0]:

$$A = \{0/9,1/9,\dots,9/9\}, |A| = 10.$$

The environment assigns a reward based on the proximity between the selected volume and a latent optimal volume level $v \in [0.3,0.9]$, drawn randomly at each episode. The reward function is defined as:

$$R(s,a) = -|a - v|$$

To incorporate user feedback, integrating real-world human-in-the-loop learning, explicit feedback is solicited with a probability of 20% at each time step. Feedback can have two types: relative (positive, negative or neutral) or direct input (preferred volume).

- Positive feedback: $R = +1.0$;
- Negative feedback: $R = -1.0$;
- Neutral feedback: $R = 0.0$;
- Scalar volume input: $R = +1.0$, and action is updated accordingly.

The agent employs the Q-learning algorithm governed by an E-greedy policy. Where the discount factor is set to 0.9, and the learning rate to 0.1. At each step, the agent selects the volume level with the highest Q-value in the current state with probability 1-E, and randomly with probability E. Training proceeds over 1000 episodes, each comprising 20-time steps.

The learning environment extends the interface from the Gymnasium library. It integrates a mechanism to inject external variables into the environment, effectively reading the context values being produced by sensors.

As shown in Figure 22, results suggest that the RL Value node applied to the Volume preference is learning over time. On the left graph, while there is some fluctuation, it is shown an increasing trend in the total reward per episode, meaning that the model is positively learning. The right graph illustrates the decline in the epsilon or exploration rate, meaning as time progresses the model explored less different volume settings.

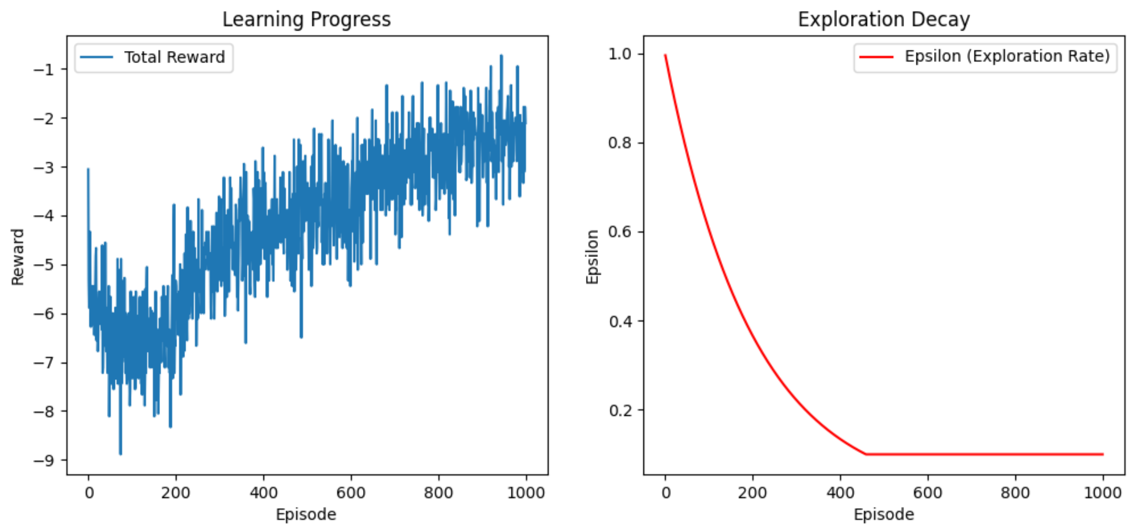

Figure 22 Total Reward and exploration decay over episodes

The proposed user structure supports real-time adaptation through mechanisms such as averaging, temporal decay, and exponentially weighted moving averages, allowing the system to reflect evolving user behavior while preserving contextual semantics.

Overall, these results confirm that the combination of the graph-based structure and reinforcement learning approach is successfully adapting to user preferences, gradually reducing unnecessary variations while improving responsiveness to user needs. The tabular approach is well-suited to low-dimensional discrete settings, such as volume control. The integration of periodic user feedback closely relates to real-world adaptive systems and supports both supervised and reinforcement paradigms.

However, current limitations include the sparse feedback mechanism, a static optimal volume per episodeand a lack of modeling for the evolution of temporal preferences. Future work should incorporate a combination of offline and online learning from historical and streaming data to improve training speeds and adaptability.

5.3 Multi-User Environments

Modeling preferences in multi-user environments introduces a layer of complexity that surpasses individual preference modeling, as it must account not only for diverse and often conflicting user goals, but also for the social, spatial, and temporal dynamics inherent to shared systems. In intelligent environments, multiple users often coexist and interact with the same digital infrastructure, yet possess distinct habits, priorities, and comfort thresholds. This case study is designed to evaluate the capability of Caravels in balancing these preferences in real-time, while maintaining perceived fairness, system efficiency, and user satisfaction. It also explores the computational modeling of preferences in such shared contexts, considering conflict resolution strategies, group decision-making frameworks, and context-aware negotiation mechanisms that reflect both individual agency and collective harmony. This case study is based on the case study presented in one publication:

- [Journal - IF 5.5 – Under review] **Rafael Silva**, Luis Gomes, Zita Vale "Distributed Computing for Intelligent Buildings: The Caravels Approach", Neurocomputing.

Firstly, User A manually initializes their preferences through the graphical interface, shown in Figure 19, allowing him to plant preferences and conditions without technical knowledge. As preferences are planted in the local graph structure, they become accessible by other services of Caravels. The user sets an initial value for the value nodes, resulting in the following structure shown in Figure 23.

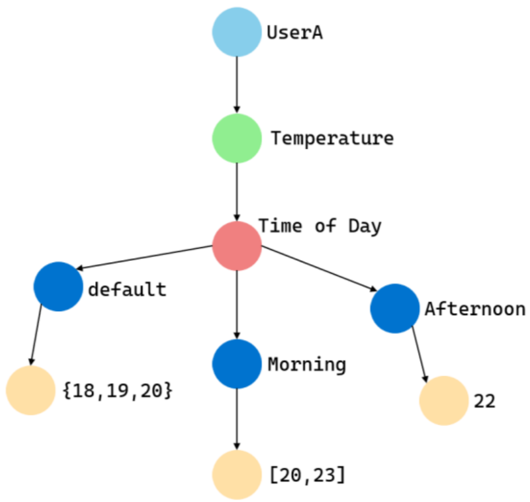

Figure 23 UserA Initial Structure

Over a period of one week, the interactions with the system are recorded. The system logs the user's actions, such as adjusting the room temperature, turning lights on/off, or opening/closing window blinds, to see if user preferences are being met, and adapts the value nodes to the historical data. The type of adaptation in the value node depends on the type. For unitary values, the value node is updated averaging the current value with the new input value. For range values, if the input is out of the interval, it automatically shifts the range to the left or to the right, depending on the side of the range the input value is on. To increase the size of the range the update of the range boundaries uses an exponentially weighted moving average with a smoothing factor of 0.1. For list type values, the value node is appended with the new input value, and each value has a temporal decay. In Figure 24, it is shown the results of these interactions, where all value nodes were updated.

To simulate real-world mobility, User B from another caravel visits the household. The system automatically imports the user's preference structure, preserving the contextual logic. This structure is similar to User A structure shown in Figure 24. When services harvest the preferences, they will detect the User B structure and behave according to his needs, allowing seamless transitions without user reconfiguration.

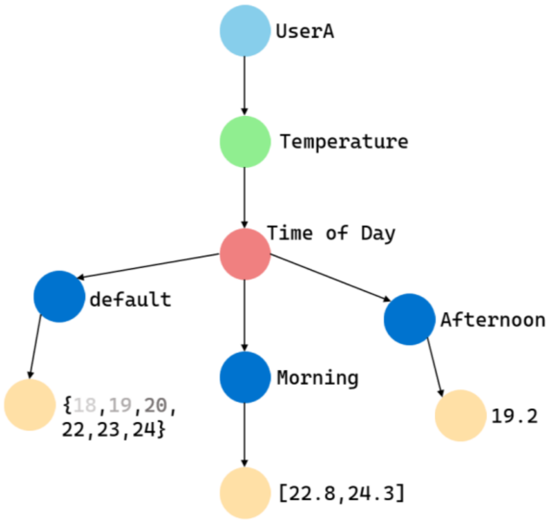

Figure 24 UserA structure after several interactions

In some hours, User A and User B are present in the same room, making the system deal with multi-user conflict. Their individual structures were maintained independently and triggered during co-occupancy. When the service tries to harvest the structures, the system first trims the structures to isolate the desired preference branch, then grafts the structures to produce a merged temporary structure. This temporary structure is shown in Figure 25, as well as the illustration of the conflict resolution. For this scenario, a dynamic weight-based strategy was chosen. This adjusts influence ratios based on occupancy time and frequency of override actions.

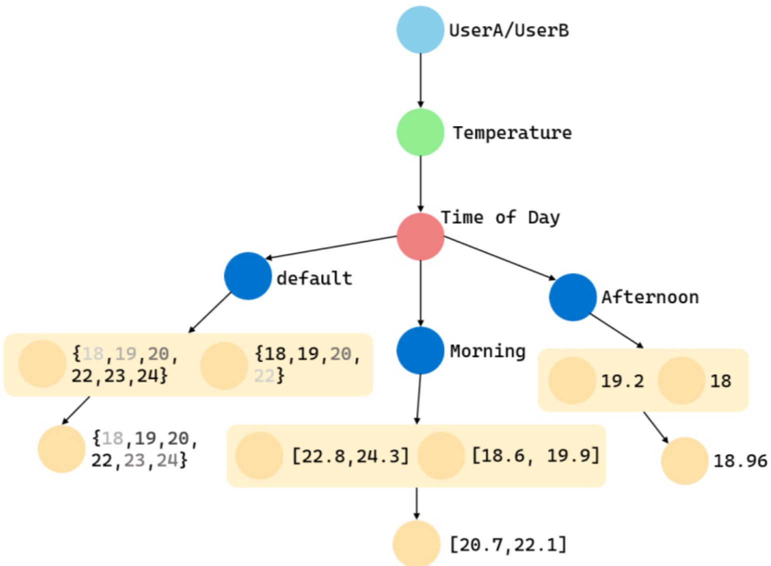

Figure 25 Temporary structure and conflict resolution illustration

Caravels uses a context-aware approach to user preference modeling in intelligent environments, which enables service abstraction from preference conflicts within shared or multi-user spaces. By leveraging a graph-based structure, Caravels enables the specification of user preferences, where conditions such as time of day or environmental parameters dynamically modulate the applicability of value nodes. These conditions can be used by multiple users, fostering the collective growth of the user structure. Moreover, the portability of user preference structures across different caravels facilitates seamless transitions between environments, upholding user agency without necessitating manual reconfiguration.

Additionally, Caravels introduces an innovative strategy for multi-user conflict resolution through the grafting of individual preference structures and the application of dynamic weight-based influence ratios. This approach respects the autonomy of individual users while enabling collaborative decision-making grounded in factors such as occupancy duration and override frequency. The system thus ensures perceived fairness and maintains user satisfaction without resorting to static prioritization schemes. This study demonstrates that Caravels not only meets the demands of real-time, socially-aware preference management but also lays a foundation for scalable, generalizable solutions across diverse modalities. Future work may explore the integration of learning-based conflict resolution and enhanced transparency mechanisms, positioning Caravels as a significant contribution to the field of intelligent user-centric systems.

5.4 Intelligent Energy Community powered by Caravels: A fully deployed case study

This case study regards the implementation of Caravels in an Energy Community, ultimately achieving an intelligent energy community. The primary objective is to evaluate Caravels ability to integrate different energy models in already deployed energy communities. This energy community has multiple energy optimization models which will be presented separately: (i) a Peer-to-peer auction-based model for transactions among members (ii) a voluntary-basis Demand Response program, and (iii) energy storage system optimization. This case study is based on the case studies presented in two publications:

- [Journal] Rita Costa, **Rafael Silva**, Ricardo Faia, Luis Gomes, Pedro Faria, Zita Vale (2024) "Empowering energy management in smart buildings: A comprehensive study on distributed energy storage systems for sustainable consumption", published in Energy and Building in 2024, doi: 10.1016/j.enbuild.2024.114953;
- [Conference] **Rafael Silva**, Luis Gomes, Zita Vale (2025) "Caravels: a Decentralized Container-Based Infrastructure for Sustainable Human-Centric Intelligent Energy Communities", presented in 2024 International Conference on Soft Computing Models in Industrial and Environmental Applications (SOCO 2024), doi: 10.1007/978-3-031-75013-7_23;
- This case study served as a foundation for GECAD demonstrations during the FCT evaluation. In this evaluation, GECAD received the EXCELLENT classification, the highest possible.

All the models discussed in this case study have been integrated within the same environment and energy community, comprising 27 members. This community includes 7 real members (2 office buildings, 5 residential buildings) and 20 simulated members. As there is no physical connection between the nodes, it is considered a virtual energy community.

In Caravels, each member owns their own computation node, which is achieved by combining both physical and virtual machines. The real members' systems are hosted on Raspberry Pi 4 devices. Since the members are physically distant from one another, this setup allows the system to test its integration across different networks. The simulated members' systems operate on virtual machines, enabling the creation of various profiles. This setup allows for the simulation of different types of buildings, such as residential, office, or industrial structures, and tests how the models behave under various circumstances. Additionally, the community node operates on a separate machine.

As seen in Figure 26, each model developed in this case study consists of two separate containers. Since these models function at the community level, one component is hosted on the community node, while the other component is deployed on each member's node. This setup allows the main component to communicate with the member's system and perform specific actions.

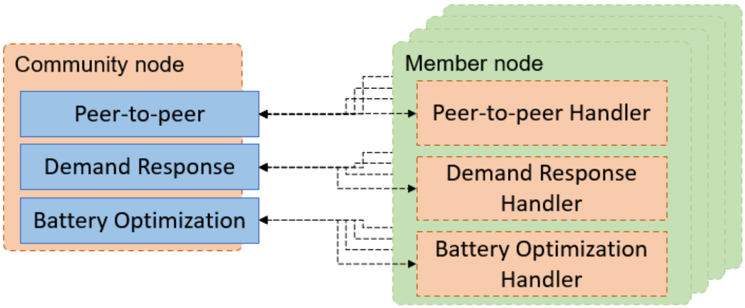

Figure 26 Communications between containers on the different nodes

As seen in Figure 27, the services hosted by the community node are exposed to the whole community. Given that the community services have active behavior, containers corresponding to the member’s handler also need to be exposed to enable the community-member communication.

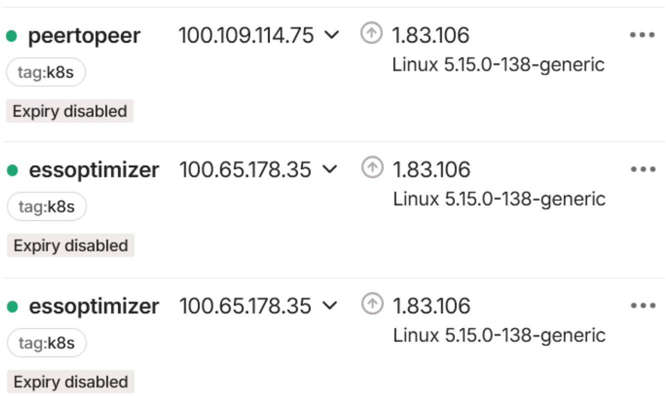

Figure 27 Tailscale interface showing exposed community services (print screen)

5.4.1 Peer-to-peer

To facilitate energy sharing among community members, an auction-based peer-to-peer (P2P) model has been integrated into Caravels for day-ahead P2P transactions. This model relies on energy forecasting tools for consumption and generation. The day-ahead forecasts provide 24 hourly energy values, which the P2P container then converts into the amounts of energy that each member intends to buy or sell in the P2P system. Figure 28 illustrates the communications between the community and the members.

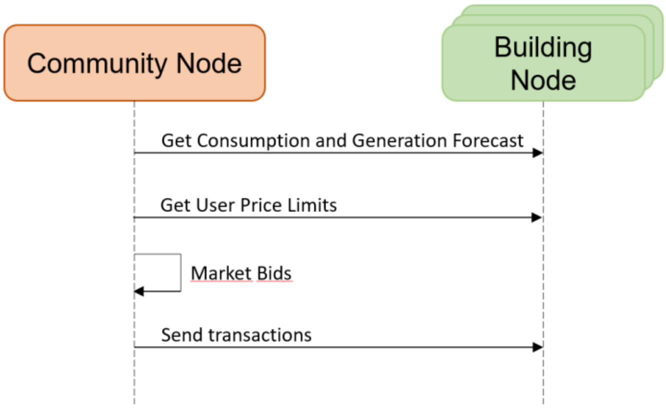

Figure 28 Peer-to-peer node communications

Before participating in the P2P auction, each community member must define the prices to which they are willing to buy or sell the energy. Each member has access to the energy market price and sets their P2P price relative to the market price. Figure 29 shows the interface where members can set their prices. The prices are defined in Euros (EUR) per kWh.

Figure 29 User interface to set energy price (print screen from the application)

To the members participating in the P2P, the community node requests the day-ahead consumption and generation forecast of each member, as well as their defined price. This makes the members submit their bids to the energy community operator through HTTP communication, detailing the energy amount in kWh and the price in EUR per kWh. After aggregating the required data of each member, the community node then starts the auction process, which involves matching bids of every member.

The P2P model allows for direct transactions between peers, and the outcomes of these transactions are subsequently communicated to each member. Participation in the auction-based P2P model is voluntary; however, once a member chooses to participate, they are obligated to fulfill the results of their transactions. For example, if a member sells 0.5 kWh, they must supply 0.5 kWh the following day, even if they need to purchase that energy from their retailer to meet that obligation.

The P2P auction is conducted on a day-ahead basis, covering all 24 hours of the following day. The P2P matrix displays all transactions among members, as seen in Figure 30. In Caravels, this information is only accessible through the community operator interface, presented in terms of energy values (kWh) or economic values (EUR). On the member side, individuals can view information related to P2P transactions, but only those transactions involving their own accounts are displayed. Transactions between other members are not publicly available to them.

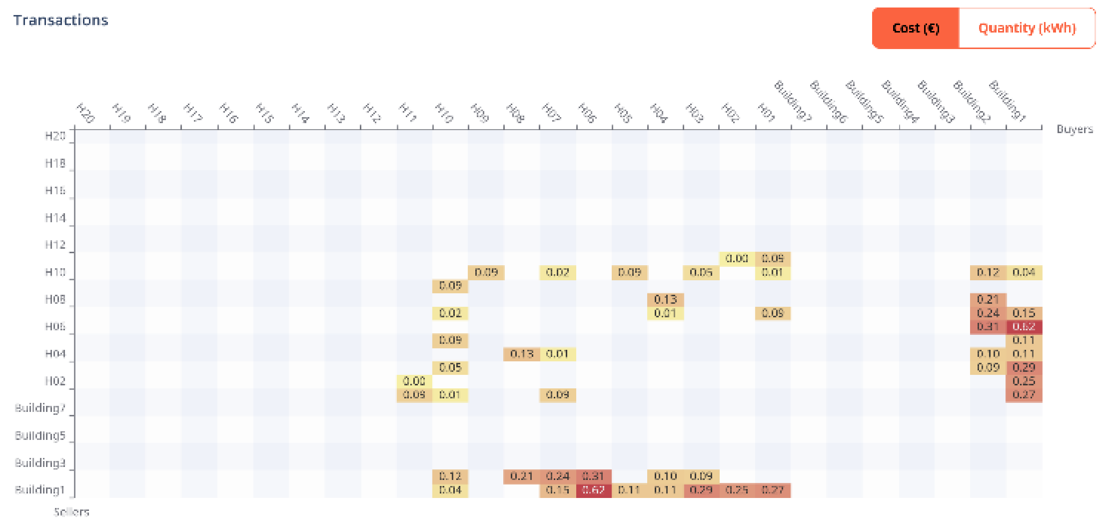

Figure 30 P2P energy transaction matrix between members (print screen from the application)

5.4.2 Demand Response

To support the energy community operator, a voluntary DR program designed for energy communities was integrated into Caravels, as published in reference (Barreto et al., 2024). This program enables the community operator to balance energy consumption and generation, allowing the community to become independent from external energy suppliers.

The integration of the DR program presents greater complexity compared to the P2P model, primarily because it involves the coordination and synchronization of multiple processes that operate over varied time frames and across different nodes in the network. As depicted in Figure 31 and similarly to the P2P model, the DR program is supported by energy forecast models deployed within the community operator and its members. Specifically, the DR program requires a day-ahead forecast of energy consumption, generation, and flexibility. The forecasts provided by each member are aggregated to calculate the general forecast of the community.

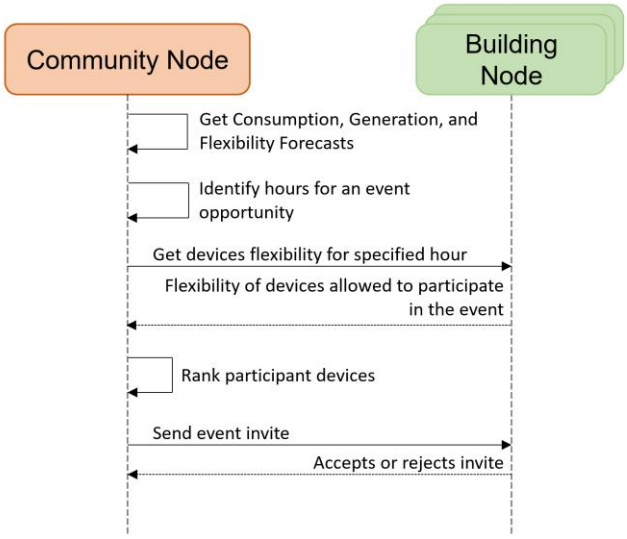

Figure 31 DR node communication

After the calculation of the community energy forecast, the DR program begins by identifying periods when DR events can be launched. A period, p, is considered viable for a DR event if the following two constraints are satisfied:

$$consumption_p > generation_p \tag{1}$$

$$consumption_{p<} - flexibility_p \leq generation_p \tag{2}$$

where consumption, generation, and flexibility represent the sum of values of the energy community (i.e., the sum of each individual member's value).

Once viable DR event periods are identified, the community operator consults its members to assess their willingness to participate and communicate the forecasted flexibility, to the determined hour of the event, of each IoT device. Members respond by indicating the flexibility they can offer for both reduction and shifting. It is important to note that only the IoT devices specified in this response will be considered for participation in the event. If a member chooses not to participate with a particular IoT device, they simply need to exclude it from their response.

The community operator aggregates the responses from all members and ranks the members according to the following four metrics:

- Total number of participations vs. flexibility provided;
- Participation rate vs. flexibility provided;

- Flexibility available for reductions in the upcoming event vs. effort rate;
- Cost of shifting vs. effort rate.

Clustering algorithms are used to group participants, and scores are assigned for each metric. After the assignment of these scores, a fairness mechanism is implemented to ensure equitable opportunities for all members, particularly for those who may not have participated in previous DR events. This will produce a list ordered by the fair score of each participant.

After the ranking process, the system selects members to invite based on their position on the list. The number of participants invited to the event will depend on the amount of energy that needs to be reduced at that time and the expected flexibility that participants can offer. This means that not all participants from the ranked list will necessarily be invited; those who are not invited will remain as reserve participants.

Once the community operator identifies which members will be invited to the DR event, they send out invitations to those members. Each member has the option to accept or decline the invitation. If a member declines, the community operator will invite the next member on the ordered scored list.

During the DR event, the community operator is responsible for monitoring the balance to ensure compliance. Throughout the one-hour event, a dedicated monitoring system checks the community's energy balance every 10 minutes. Every 10 minutes, if consumption surpasses generation, the community operator can invite additional participants in real-time.

The successful integration of the DR program demonstrated Caravels' capacity to manage multiple processes across various communications, announcements, and invitations. Figure 32 shows the DR event monitoring chart where the energy balance of the community was tracked during the one-hour event. It is evident that the energy balance was above zero during two periods, specifically at 40 and 50 minutes after the event began. During these times, additional participants who had initially remained as reserves were invited to join the event, resulting in a decrease in the community's energy balance. The corrections are indicated by their designated intervals: correction 4 occurs at 40 minutes, and correction 5 at 50 minutes.

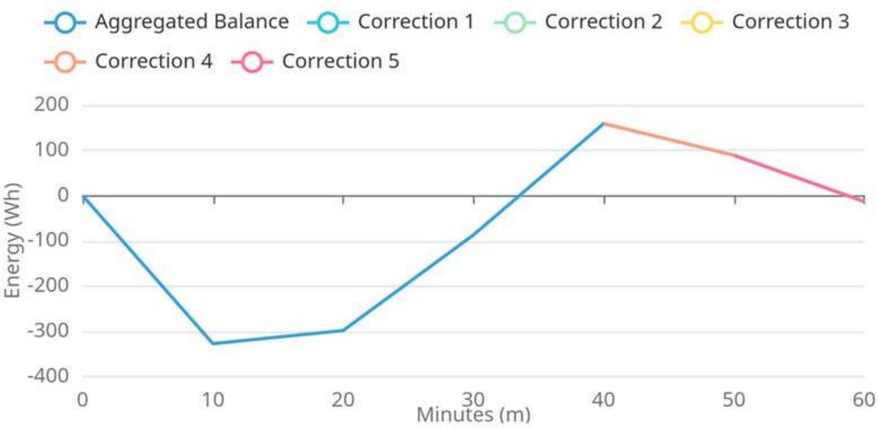

Figure 32 DR event monitoring (print screen from the application)

The execution of the DR program within the energy community was made possible through the integration of energy forecast models, specifically for consumption, generation, and flexibility. Additionally, the entire Kubernetes infrastructure was tested during this event, as communication among members was crucial for the success of the DR event.

5.4.3 Energy Storage Management

An Energy Storage System (ESS) model is implemented to manage a set of batteries deployed within an energy community. These batteries are owned by the community's members but are centrally managed by a community operator. The deterministic ESS management model schedules the charging and discharging of each battery unit, taking into account community demand and energy prices. The model integrated was published in (Costa et al., 2024).

In this case study, there are six batteries distributed between two members. One member has three batteries: two with a capacity of 3.6 kWh each and one with a capacity of 2.5 kWh. The other member also has three batteries: two with a capacity of 2.5 kWh each and one with a capacity of 3.6 kWh.

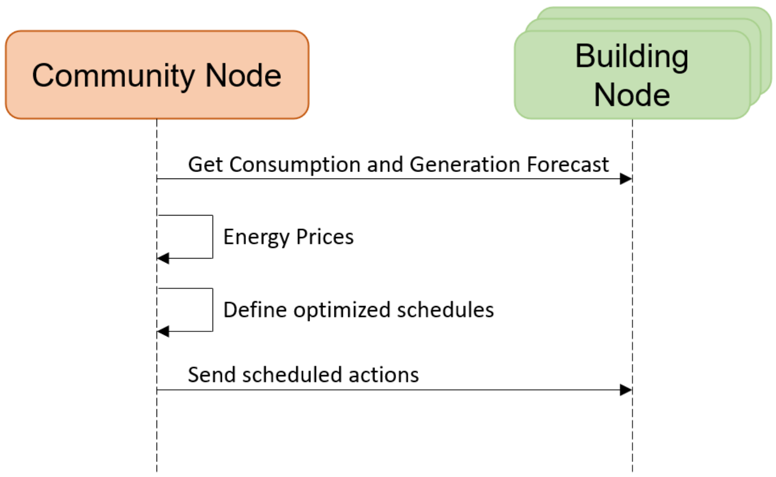

Figure 33 Energy storage optimization sequence diagram

To optimize the battery schedules, the model utilizes energy forecast models for both consumption and generation on a day-ahead basis. It generates a schedule for the next 24 hours, with the charging and discharging actions. This schedule is communicated to the member-owner of the battery, who will ensure that the planned schedule is executed as intended. The control of the ESS is done using the HTTP protocol available in the battery inverter. In Caravels, the ESS is modeled as an IoT device.

Figure 34 illustrates the state of charge of the batteries, both individually for past periods and for future periods scheduled by the optimizer. In this case study, only the planned charging and discharging actions were communicated to the batteries.

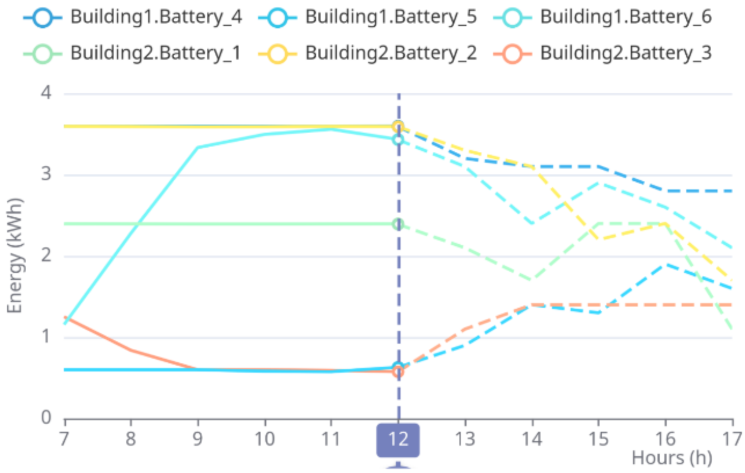

Figure 34 State of charge profile of ESS considering past periods and planned actions (print screen from the application)

This implementation of the model allows for optimized centralized scheduling of ESS at the community level. To enable this centralized management, Caravels utilized the data and service-sharing resource functionality, already mentioned before. This feature enables the owner of a resource to share both the data provided by that resource and the service controlling it with another member, in this case, the community operator. As a result, the batteries located in the office buildings and shared by both members with the community operator take into account not only the needs of their individual members but also the overall context of the community.

6 Conclusion

This chapter will focus on the major conclusions drawn from this dissertation. Additionally, the objectives achieved will be described. Several functionalities and capabilities are proposed for future integration in order to contribute to the Caravels' evolution.

6.1 Main Conclusions

With the completion of this work, it is possible to conclude that all objectives have been met and all research questions have been addressed. Throughout the project's development, unexpected obstacles and difficulties emerged that required re-planning of certain options in order to achieve optimal results. Nevertheless, the final outcomes align with the initially defined objectives.

The first objective O1 ("To design and develop a multi-domain solution for context-aware intelligent community management supported by IoT devices in retrofitted buildings") and corresponding research question, RQ1 ("Can a multi-domain solution for intelligent community management be developed and implemented using IoT devices, artificial intelligence, secure information sharing, and distributed computing?"), were successfully addressed through the design and implementation of a distributed, container-based architecture for intelligent community management. The architecture, presented in section 4.1, integrates IoT devices across buildings and public infrastructure, enabling location-aware interactions and data acquisition. The system demonstrates a clear ability to adapt to each building's scenario and manage diverse subsystems, as seen in the case studies in section 5, such as energy, preferences, and environmental monitoring, through coordinated, scalable services. The deployment of distributed entities and middleware effectively validated the feasibility of combining AI, secure information exchange, and distributed computation to support real-time, context-sensitive decision-making in community environments.

Second objective O2 ("To investigate and implement artificial intelligence techniques, such as machine learning and optimization algorithms, for optimizing community resources considering the users preferences and needs") and research question RQ2 ("Can artificial intelligence models be effectively utilized to optimize resource usage, share useful data, consider user needs, improve service mobility, and enhance social engagement within citizen communities?") were addressed by integrating machine learning and optimization techniques into the system's control and decision-making layers, using the mechanism described in sections 4.2.1, 4.2.2, and 4.2.3, namely in the integration of machine learning, P2P, DR, and energy storage optimization. These models were demonstrated in the case studies of sections 5.2, and 5.4 and used techniques such as predictive modeling, adaptive scheduling, and multi-objective optimization to reduce energy waste, anticipate peak loads and to reduce energy costs. In parallel, social engagement was fostered through AI-driven mechanisms that operate at community-level activities, and through user-friendly interfaces.

As demonstrated in 5.2 and 5.3 sections, in response to O3 ("To explore the development of dynamic and adaptive user modeling systems to enhance user interaction and provide personalized support to community members") and RQ3 ("Can dynamic and adaptive user modeling be designed to provide personalized and user-friendly interactions for managing citizen communities?"), a dynamic and adaptive user modeling framework was developed to support personalized system behavior across multiple domains. The system captures evolving user preferences, contextual cues, and behavioral patterns, allowing for longitudinal adaptation and tailored interaction strategies. This modeling capability supports context-aware system behavior, enhancing the accessibility and usability of community-level and building-level services.

The system takes into account the incorporation of strong data governance principles aligned with the GDPR, related to O4 ("To develop a container-based distributed architecture to support the seamless integration of multiple management models in intelligent communities") and RQ4 ("What technological tools could be applied to create a secure and distributed infrastructure for containers in heterogeneous machines?"). Section 4.1 presented a container-based distributed architecture enabled by the use of Kubernetes and Docker containers. This infrastructure used state of the art technologies to provide the basis for the Caravels solution. Secure communication protocols and access controls were integrated throughout the system architecture. These measures ensured that all personal data was processed transparently and securely, enabling trust in the collaborative sharing mechanisms fundamental to intelligent community operation. The measures presented in section 3.2 were implemented in each service throughout all case studies.

The fifth objective O5 ("To test and evaluate the feasibility, effectiveness, efficiency and impact of the proposed solution through real-world case studies and user testing with a diverse range of community members") was addressed in section 5 through empirical validation of the proposed system in real-world scenarios. Pilot deployments and user testing demonstrated the system's effectiveness in managing complex environments, its efficiency in optimizing resource consumption, and its positive impact on user engagement and satisfaction. The evaluations provided evidence that the architecture is both technically viable and socially inclusive, reinforcing its relevance for real-world intelligent community applications.

Intelligent communities management solutions hold the potential to revolutionize the way we manage our communities and achieve sustainability, efficiency, and overall well-being. By breaking down silos, embracing flexibility and adaptability, prioritizing user-friendliness, and data sharing while preserving security, we can create intelligent community management systems that empower communities to thrive.

This dissertation resulted in 7 scientific publications, of which 2 have been published in journals and the rest have been presented at scientific conferences. These papers prove that the proposed solution is innovative and fills the identified gaps in the state of the art. Moreover, the research presented in this work has not remained in theoretical development; it has led to real-world applications through pilot implementations and case studies.

This dissertation has also served as a foundational platform for academic continuity and dissemination. Several research avenues emerging from this work have been taken up by bachelor's theses, allowing students to explore specific components of the system, such as distributed optimization, data privacy mechanisms, and IoT integration. These derivative projects reflect the extensibility of the proposed architecture.

Furthermore, new research domains have emerged from this work, including cross-building cooperation, adaptive human-centric control systems, and preference-aware AI governance. These domains open promising lines of inquiry for future interdisciplinary research that bridges artificial intelligence, urban informatics, human-computer interaction, and data ethics. In this sense, the dissertation not only advances knowledge but also establishes fertile groundwork for ongoing innovation and application in intelligent community systems.

6.2 Future work

The proposed work has demonstrated its value in achieving all initially defined objectives across various contexts. However, there are some adjustments and upgrades that can be made to enhance the potential uses of Caravels.

The implementation of a distributed architecture for managing intelligent community systems demonstrated robust performance, scalability, and resilience across multiple subsystems. By decentralizing decision-making and enabling localized optimization, the architecture effectively reduced bottlenecks and single points of failure, while promoting modularity and inter-building coordination. However, the principal drawback identified was the increased complexity in maintaining and developing a high number of services. This challenge underscores the need for further research into effective management in such modular environments.

In the domain of user preference modeling, the integration of user feedback and contextual information allowed for a more nuanced and personalized control strategy across services. Nevertheless, current implementations relied on relatively simple structures. A promising direction for future work lies in the application of Graph Neural Networks (GNNs), which can natively encode and process structured relational data, such as user-to-preference and context-to-behavior graphs. GNNs offer the potential to enhance preference inference, model non-linear dependencies, and adapt dynamically to evolving user patterns, thereby improving both personalization and system efficiency at scale.

Future work will explore the incorporation of purpose-designed autonomous agents tasked with achieving explicit objectives, such as minimizing energy consumption, maximizing self-consumption, by constructing and evolving graph-based structures analogous to those used in the user preference module, dynamically adapting service harvesting in real time.

References

Adel, A., & HS Alani, N. (2024). Human-Centric Collaboration and Industry 5.0 Framework in Smart Cities and Communities: Fostering Sustainable Development Goals 3, 4, 9, and 11 in Society 5.0. *Smart Cities 2024, Vol. 7, Pages 1723-1775*, *7*(4), 1723–1775. https://doi.org/10.3390/SMARTCITIES7040068

Aguilar, J., Garces-Jimenez, A., R-Moreno, M. D., & García, R. (2021). A systematic literature review on the use of artificial intelligence in energy self-management in smart buildings. *Renewable and Sustainable Energy Reviews*, *151*, 111530. https://doi.org/10.1016/J.RSER.2021.111530

Al Nuaimi, E., Al Neyadi, H., Mohamed, N., & Al-Jaroodi, J. (2015). Applications of big data to smart cities. *Journal of Internet Services and Applications*, *6*(1), 1–15. https://doi.org/10.1186/S13174-015-0041-5/FIGURES/3

Aljawarneh, M., Zullqar, S., Bhutto, A., Das Dhomeja, L., Malkani, Y. A., Arafat Malkani, Y., & Shah, S. M. (2021). *Resolving User Conflicts in Multi-user Context-aware Home Environment*. https://doi.org/10.21203/RS.3.RS-1045928/V1

Anawar, M. R., Wang, S., Azam Zia, M., Jadoon, A. K., Akram, U., & Raza, S. (2018). Fog Computing: An Overview of Big IoT Data Analytics. *Wireless Communications and Mobile Computing*, *2018*(1), 7157192. https://doi.org/10.1155/2018/7157192

Ardabili, S. F., Mosavi, A., Ghamisi, P., Ferdinand, F., Varkonyi-Koczy, A. R., Reuter, U., Rabczuk, T., & Atkinson, P. M. (2020). COVID-19 Outbreak Prediction with Machine Learning. *Algorithms 2020, Vol. 13, Page 249*, *13*(10), 249. https://doi.org/10.3390/A13100249

Atas, M., Felfernig, A., Polat-Erdeniz, S., Popescu, A., Tran, T. N. T., & Uta, M. (2021). Towards psychology-aware preference construction in recommender systems: Overview and research issues. *Journal of Intelligent Information Systems*, *57*(3), 467–489. https://doi.org/10.1007/S10844-021-00674-5/TABLES/5

Auda, J., Weber, D., Voit, A., & Schneegass, S. (2018). Understanding user preferences towards rule-based notification deferral. *Conference on Human Factors in Computing Systems - Proceedings*, *2018-April*. https://doi.org/10.1145/3170427.3188688

Ayoub Shaikh, T., Rasool, T., & Rasheed Lone, F. (2022). Towards leveraging the role of machine learning and artificial intelligence in precision agriculture and smart farming. *Computers and Electronics in Agriculture*, *198*, 107119. https://doi.org/10.1016/J.COMPAG.2022.107119

Barone, G., Buonomano, A., Forzano, C., Palombo, A., & Russo, G. (2023). The role of energy communities in electricity grid balancing: A flexible tool for smart grid power distribution optimization. *Renewable and Sustainable Energy Reviews*, *187*, 113742. https://doi.org/10.1016/J.RSER.2023.113742

Barreto, R., Gomes, L., & Vale, Z. (2024). Smart Solution for Energy Communities: Integrating Demand Response and Unsupervised Learning Evaluation Metrics. *2024 IEEE 22nd Mediterranean Electrotechnical Conference, MELECON 2024*, 144–149. https://doi.org/10.1109/MELECON56669.2024.10608742

Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics*, *92*(8–9), 1787–1794. https://doi.org/10.1016/J.JPUBECO.2008.04.010

Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M., & Lee, M. (2009). Natural actor–critic algorithms. *Automatica*, *45*(11), 2471–2482. https://doi.org/10.1016/J.AUTOMATICA.2009.07.008

Bibri, S. E., Huang, J., Jagatheesaperumal, S. K., & Krogstie, J. (2024). The synergistic interplay of artificial intelligence and digital twin in environmentally planning sustainable smart cities: A comprehensive systematic review. *Environmental Science and Ecotechnology*, *20*, 100433. https://doi.org/10.1016/J.ESE.2024.100433

Bohn, J., Coroamă, V., Langheinrich, M., Mattern, F., & Rohs, M. (2004). Living in a world of smart everyday objects - Social, economic, and ethical implications. *Human and Ecological Risk Assessment*, *10*(5), 763–785. https://doi.org/10.1080/10807030490513793

Bousbiat, H., Bousselidj, R., Himeur, Y., Amira, A., Bensaali, F., Fadli, F., Mansoor, W., & Elmenreich, W. (2023). *Crossing Roads of Federated Learning and Smart Grids: Overview, Challenges, and Perspectives*. https://arxiv.org/abs/2304.08602v1

Butt, U. J., Davelis, A., Abbod, M., Eghan, C., & Agbo, H. M. (2022). Improving Learning Experience and Privacy in Education Using the Power of Big Data and Artificial Intelligence. *Integrated Business Models in the Digital Age: Principles and Practices of Technology Empowered Strategies*, 371–424. https://doi.org/10.1007/978-3-030-97877-8_11

Caiazza, C., Giordano, S., Luconi, V., & Vecchio, A. (2021). Edge Computing vs Centralized Cloud: Impact of Communication Latency on the Energy Consumption of LTE Terminal Nodes. *Computer Communications*, *194*, 213–225. https://doi.org/10.1016/j.comcom.2022.07.026

Castro, C., & Lopes, C. (2022). Digital Government and Sustainable Development. *Journal of the Knowledge Economy*, *13*(2), 880–903. https://doi.org/10.1007/S13132-021-00749-2/TABLES/6

Cespedes-Cubides, A. S., & Jradi, M. (2024). A review of building digital twins to improve energy efficiency in the building operational stage. *Energy Informatics*, *7*(1), 1–31. https://doi.org/10.1186/S42162-024-00313-7/FIGURES/7

Chang, V. (2021). An ethical framework for big data and smart cities. *Technological Forecasting and Social Change*, *165*, 120559. https://doi.org/10.1016/J.TECHFORE.2020.120559

Chang, Z., Liu, S., Xiong, X., Cai, Z., & Tu, G. (2021). A Survey of Recent Advances in Edge-Computing-Powered Artificial Intelligence of Things. *IEEE Internet of Things Journal*, *8*(18), 13849–13875. https://doi.org/10.1109/JIOT.2021.3088875

Chen, X., Witczak-Krempa, W., Faulkner, T., Norris, J. P., Gehrels, N., Scargle, J. D., Zeng, Q.-B., Chen, S., Lü, R., Pascual, R., Pascual, J., Ortiz, J., & Salom, J. (2020). Review and comparison of methodologies and tools to perform energy efficiency projects in public buildings in MED areas. *IOP Conference Series: Earth and Environmental Science*, *410*(1), 012022. https://doi.org/10.1088/1755-1315/410/1/012022

Chen, Z., Sivaparthipan, C. B., & Muthu, B. A. (2022). IoT based smart and intelligent smart city energy optimization. *Sustainable Energy Technologies and Assessments*, *49*, 101724. https://doi.org/10.1016/J.SETA.2021.101724

Cheng, Y., Xie, Q., & Ma, Y. (2020). A Research of Instant Survey Based on Typical Scenario. *IOP Conference Series: Materials Science and Engineering*, *790*(1), 012003. https://doi.org/10.1088/1757-899X/790/1/012003

Ciasullo, M. V., Troisi, O., Grimaldi, M., & Leone, D. (2020). Multi-level governance for sustainable innovation in smart communities: an ecosystems approach. *International Entrepreneurship and Management Journal*, *16*(4), 1167–1195. https://doi.org/10.1007/S11365-020-00641-6/FIGURES/5

Clements-Croome, T. D. J. (1997). What do we mean by intelligent buildings? *Automation in Construction*, *6*(5–6), 395–400. https://doi.org/10.1016/S0926-5805(97)00018-6

Coignard, J., Rigo-Mariani, R., & Debusschere, V. (2023). Individual and collective objectives in an energy community with network constraints. *Sustainable Cities and Society*, 105083. https://doi.org/10.1016/J.SCS.2023.105083

Cosentino, G., & Giannakos, M. (2023). Multisensory Interaction and Analytics to Enhance Smart Learning Environments: A Systematic Literature Review. *IEEE Transactions on Learning Technologies*, *16*(3), 414–430. https://doi.org/10.1109/TLT.2023.3243210

Costa, R., Silva, R., Faia, R., Gomes, L., Faria, P., & Vale, Z. (2024). Empowering energy management in smart buildings: A comprehensive study on distributed energy storage systems for Sustainable consumption. *Energy and Buildings*, *324*, 114953. https://doi.org/10.1016/J.ENBUILD.2024.114953

Dasgupta, A., Handosa, M., Manuel, M., & Gračanin, D. (2019). A user-centric design framework for smart built environments: A mixed reality perspective. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *11587 LNCS*, 124–143. https://doi.org/10.1007/978-3-030-21935-2_11

De Guimarães, J. C. F., Severo, E. A., Felix Júnior, L. A., Da Costa, W. P. L. B., & Salmoria, F. T. (2020). Governance and quality of life in smart cities: Towards sustainable development goals. *Journal of Cleaner Production*, *253*, 119926. https://doi.org/10.1016/J.JCLEPRO.2019.119926

de São José, D., Faria, P., & Vale, Z. (2021). Smart energy community: A systematic review with metanalysis. *Energy Strategy Reviews*, *36*, 100678. https://doi.org/10.1016/J.ESR.2021.100678

Dietrich, F., & List, C. (2013). Where do preferences come from? *International Journal of Game Theory*, *42*(3), 613–637. https://doi.org/10.1007/S00182-012-0333-Y/METRICS

Ding, Y., Han, S., Tian, Z., Yao, J., Chen, W., & Zhang, Q. (2021). Review on occupancy detection and prediction in building simulation. *Building Simulation 2021 15:3*, *15*(3), 333–356. https://doi.org/10.1007/S12273-021-0813-8

DIRECTIVE (EU) 2018/2001. (2018). In *Official Journal of the European Union*.

DIRECTIVE (EU) 2019/944 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL. (n.d.). Retrieved February 7, 2025, from https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32019L0944

Djehaiche, R., Aidel, S., Sawalmeh, A., Saeed, N., & Alenezi, A. H. (2023). Adaptive Control of IoT/M2M Devices in Smart Buildings Using Heterogeneous Wireless Networks. *IEEE Sensors Journal*, *23*(7), 7836–7849. https://doi.org/10.1109/JSEN.2023.3247007

Dwivedi, D., Babu, K. V. S. M., Yemula, P. K., Chakraborty, P., & Pal, M. (2022). *Evaluation of Energy Resilience and Cost Benefit in Microgrid with Peer-to-Peer Energy Trading*. https://arxiv.org/abs/2212.02318v1

Eke, C. I., Norman, A. A., Shuib, L., & Nweke, H. F. (2019). A Survey of User Profiling: State-of-the-Art, Challenges, and Solutions. *IEEE Access*, *7*, 144907–144924. https://doi.org/10.1109/ACCESS.2019.2944243

El Naqa, I., & Murphy, M. J. (2015). What Is Machine Learning? *Machine Learning in Radiation Oncology*, 3–11. https://doi.org/10.1007/978-3-319-18305-3_1

Elalfy, D. A., Gouda, E., Kotb, M. F., Bureš, V., & Sedhom, B. E. (2024). Comprehensive review of energy storage systems technologies, objectives, challenges, and future trends. *Energy Strategy Reviews*, *54*, 101482. https://doi.org/10.1016/J.ESR.2024.101482

Emami-Naeini, P., Breda, J., Dai, W., Kohno, T., Laine, K., Patel, S., & Roesner, F. (n.d.). *Understanding People's Con-cerns and Attitudes Toward Smart Cities*. *23*, 24. https://doi.org/10.1145/3544548.3581558

Energy communities. (n.d.). Retrieved December 12, 2023, from https://energy.ec.europa.eu/topics/markets-and-consumers/energy-communities_en

Energy Communities Repository - General information - European Commission. (n.d.). Retrieved December 12, 2023, from https://energy-communities-repository.ec.europa.eu/energy-communities-repository-energy-communities/energy-communities-repository-general-information_en

Energy saving: EU action to reduce energy consumption | News | European Parliament. (n.d.). Retrieved December 12, 2023, from https://www.europarl.europa.eu/news/en/headlines/society/20221128STO58002/energy-saving-eu-action-to-reduce-energy-consumption

Engelhardt, J., Zepter, J. M., Gabderakhmanova, T., & Marinelli, M. (2022). Energy management of a multi-battery system for renewable-based high power EV charging. *ETransportation*, *14*, 100198. https://doi.org/10.1016/J.ETRAN.2022.100198

Espe, E., Potdar, V., & Chang, E. (2018). Prosumer Communities and Relationships in Smart Grids: A Literature Review, Evolution and Future Directions. *Energies 2018, Vol. 11, Page 2528*, *11*(10), 2528. https://doi.org/10.3390/EN11102528

Esrafilian-Najafabadi, M., & Haghighat, F. (2021). Occupancy-based HVAC control systems in buildings: A state-of-the-art review. *Building and Environment*, *197*, 107810. https://doi.org/10.1016/J.BUILDENV.2021.107810

Fan, C., Chen, M., Wang, X., Wang, J., & Huang, B. (2021). A Review on Data Preprocessing Techniques Toward Efficient and Reliable Knowledge Discovery From Building Operational Data. *Frontiers in Energy Research*, *9*, 652801. https://doi.org/10.3389/FENRG.2021.652801/BIBTEX

Figueiredo Prudencio, R., Maximo, M. R. O. A., & Colombini, E. L. (2024). A Survey on Offline Reinforcement Learning: Taxonomy, Review, and Open Problems. *IEEE Transactions on Neural Networks and Learning Systems*, *35*(8), 10237–10257. https://doi.org/10.1109/TNNLS.2023.3250269

Fister, I., Yang, X. S., Brest, J., & Fister, D. (2013). A Brief Review of Nature-Inspired Algorithms for Optimization. *Elektrotehniski Vestnik/Electrotechnical Review*, *80*(3), 116–122. https://arxiv.org/abs/1307.4186v1

Freund, S., & Schmitz, G. (2021). Implementation of model predictive control in a large-sized, low-energy office building. *Building and Environment*, *197*, 107830. https://doi.org/10.1016/J.BUILDENV.2021.107830

Frieden, D., Roberts, J., & Gubina, A. F. (2019). Overview of emerging regulatory frameworks on collective self-consumption and energy communities in Europe. *International Conference on the European Energy Market, EEM*, *2019-September*. https://doi.org/10.1109/EEM.2019.8916222

Furszyfer Del Rio, D. D. (2022). Smart but unfriendly: Connected home products as enablers of conflict. *Technology in Society*, *68*, 101808. https://doi.org/10.1016/J.TECHSOC.2021.101808

Gao, M., Liu, K., & Wu, Z. (2010). Personalisation in web computing and informatics: Theories, techniques, applications, and future research. *Information Systems Frontiers*, *12*(5), 607–629. https://doi.org/10.1007/S10796-009-9199-3/TABLES/4

García-Monge, M., Zalba, B., Casas, R., Cano, E., Guillén-Lambea, S., López-Mesa, B., & Martínez, I. (2023a). Is IoT monitoring key to improve building energy efficiency? Case study of a smart campus in Spain. *Energy and Buildings*, *285*, 112882. https://doi.org/10.1016/J.ENBUILD.2023.112882

García-Monge, M., Zalba, B., Casas, R., Cano, E., Guillén-Lambea, S., López-Mesa, B., & Martínez, I. (2023b). Is IoT monitoring key to improve building energy efficiency? Case study of a smart campus in Spain. *Energy and Buildings*, *285*, 112882. https://doi.org/10.1016/J.ENBUILD.2023.112882

General Data Protection Regulation (GDPR). (n.d.). Retrieved December 11, 2023, from https://gdpr-info.eu/

Ghazy, N., Abdelkader, A., Zaki, M. S., & Eldahshan, K. A. (2023). An ameliorated Round Robin algorithm in the cloud computing for task scheduling. *Bulletin of Electrical Engineering and Informatics*, *12*(2), 1103–1114. https://doi.org/10.11591/EEI.V12I2.4524

Giordano, A., Ghiani, E., Pilo, F., & Rosetti, L. (2019). Planning of Energy Production and Management of Energy Resources in Local Energy Communities: The Case of Berchidda Municipality (Italy). *Proceedings 2019, Vol. 20, Page 16*, *20*(1), 16. https://doi.org/10.3390/PROCEEDINGS2019020016

Gjorgievski, V. Z., Cundeva, S., & Georghiou, G. E. (2021). Social arrangements, technical designs and impacts of energy communities: A review. *Renewable Energy*, *169*, 1138–1156. https://doi.org/10.1016/J.RENENE.2021.01.078

Gracias, J. S., Parnell, G. S., Specking, E., Pohl, E. A., & Buchanan, R. (2023). Smart Cities—A Structured Literature Review. *Smart Cities 2023, Vol. 6, Pages 1719-1743*, *6*(4), 1719–1743. https://doi.org/10.3390/SMARTCITIES6040080

Gronauer, S., & Diepold, K. (2022). Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review*, *55*(2), 895–943. https://doi.org/10.1007/S10462-021-09996-W/TABLES/4

Guo, Z., & Wang, H. (2021). A Deep Graph Neural Network-Based Mechanism for Social Recommendations. *IEEE Transactions on Industrial Informatics*, *17*(4), 2776–2783. https://doi.org/10.1109/TII.2020.2986316

Gupta, I., Singh, A. K., Lee, C. N., & Buyya, R. (2022). Secure Data Storage and Sharing Techniques for Data Protection in Cloud Environments: A Systematic Review, Analysis, and Future Directions. *IEEE Access*, *10*, 71247–71277. https://doi.org/10.1109/ACCESS.2022.3188110

Hashempour, N., Taherkhani, R., & Mahdikhani, M. (2019). *Energy performance optimization of existing buildings: A literature review*. https://doi.org/10.1016/j.scs.2019.101967

Heaslip, E., Costello, G. J., & Lohan, J. (2016). Assessing Good-Practice Frameworks for the Development of Sustainable Energy Communities in Europe: Lessons from Denmark and Ireland. *[Journal of Sustainable Development of Energy, Water and Environment Systems]*, *[4]*([3]), [307]-[319]. https://doi.org/10.13044/J.SDEWES.2016.04.0024

Hernandez, A. A., Yip, S. Y., Goh, K. M., Caballero, A. R., & Victorino, M. L. (2025). Peer-to-peer Energy Resource Sharing in Rural Communities: Enabling Technologies, Applications, and Challenges. *IEEE Access*. https://doi.org/10.1109/ACCESS.2025.3561701

Himeur, Y., Elnour, M., Fadli, F., Meskin, N., Petri, I., Rezgui, Y., Bensaali, F., & Amira, A. (2023). AI-big data analytics for building automation and management systems: a survey, actual challenges and future perspectives. *Artificial Intelligence Review*, *56*(6), 4929–5021. https://doi.org/10.1007/S10462-022-10286-2/FIGURES/12

Hopgood, A. A. (2005a). The State of Artificial Intelligence. *Advances in Computers*, *65*, 1–75. https://doi.org/10.1016/S0065-2458(05)65001-2

Hopgood, A. A. (2005b). The State of Artificial Intelligence. *Advances in Computers*, *65*, 1–75. https://doi.org/10.1016/S0065-2458(05)65001-2

Howard, J. (2019). Artificial intelligence: Implications for the future of work. *American Journal of Industrial Medicine*, *62*(11), 917–926. https://doi.org/10.1002/AJIM.23037

Hubbs, C. D., Perez, H. D., Sarwar, O., Sahinidis, N. V., Grossmann, I. E., & Wassick, J. M. (2020). *OR-Gym: A Reinforcement Learning Library for Operations Research Problems*. https://arxiv.org/abs/2008.06319v2

Ibaseta, D., García, A., Álvarez, M., Garzón, B., Díez, F., Coca, P., Pero, C. Del, & Molleda, J. (2021a). Monitoring and control of energy consumption in buildings using WoT: A novel approach for smart retrofit. *Sustainable Cities and Society*, *65*, 102637. https://doi.org/10.1016/J.SCS.2020.102637

Ibaseta, D., García, A., Álvarez, M., Garzón, B., Díez, F., Coca, P., Pero, C. Del, & Molleda, J. (2021b). Monitoring and control of energy consumption in buildings using WoT: A novel approach for smart retrofit. *Sustainable Cities and Society*, *65*, 102637. https://doi.org/10.1016/J.SCS.2020.102637

Ismagilova, E., Hughes, L., Rana, N. P., & Dwivedi, Y. K. (2020). Security, Privacy and Risks Within Smart Cities: Literature Review and Development of a Smart City Interaction Framework.

Information Systems Frontiers 2020 24:2, *24*(2), 393–414. https://doi.org/10.1007/S10796-020-10044-1

Jakobsen, K., Mikalsen, M., & Lilleng, G. (2023). A literature review of smart technology domains with implications for research on smart rural communities. *Technology in Society*, *75*, 160–791. https://doi.org/10.1016/j.techsoc.2023.102397

Jiang, T., Sun, Z., Fu, S., & Lv, Y. (2024). Human-AI interaction research agenda: A user-centered perspective. *Data and Information Management*, *8*(4), 100078. https://doi.org/10.1016/J.DIM.2024.100078

Jiang, Y., Li, X., Luo, H., Yin, S., & Kaynak, O. (2022). Quo vadis artificial intelligence? *Discover Artificial Intelligence 2022 2:1*, *2*(1), 1–19. https://doi.org/10.1007/S44163-022-00022-8

John Dian, F., Vahidnia, R., & Rahmati, A. (2020). Wearables and the Internet of Things (IoT), Applications, Opportunities, and Challenges: A Survey. *IEEE Access*, *8*, 69200–69211. https://doi.org/10.1109/ACCESS.2020.2986329

Johnston, V., Black, M., Wallace, J., Mulvenna, M., & Bond, R. (2019). A framework for the development of a dynamic adaptive intelligent user interface to enhance the user experience. *ECCE 2019 - Proceedings of the 31st European Conference on Cognitive Ergonomics: "'Design for Cognition,'"* 32–35. https://doi.org/10.1145/3335082.3335125

Kameda, T., Toyokawa, W., & Tindale, R. S. (2022). Information aggregation and collective intelligence beyond the wisdom of crowds. *Nature Reviews Psychology 2022 1:6*, *1*(6), 345–357. https://doi.org/10.1038/s44159-022-00054-y

Kanoje, S., Girase, S., & Mukhopadhyay, D. (2015). User Profiling Trends, Techniques and Applications. *International Journal of Advance Foundation and Research in Computer (IJAFRC)*, *1*(1). https://arxiv.org/abs/1503.07474v1

Kielar, P. M., & Borrmann, A. (2016). Modeling pedestrians' interest in locations: A concept to improve simulations of pedestrian destination choice. *Simulation Modelling Practice and Theory*, *61*, 47–62. https://doi.org/10.1016/J.SIMPAT.2015.11.003

Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, *79*(1), 1–14. https://doi.org/10.1007/S10708-013-9516-8/FIGURES/3

Komninos, N. (2006). The architecture of intelligent cities: Integrating human, collective and artificial intelligence to enhance knowledge and innovation. *IET Conference Publications*, *518*, 13–20. https://doi.org/10.1049/CP:20060620

Kourou, K., Exarchos, T. P., Exarchos, K. P., Karamouzis, M. V., & Fotiadis, D. I. (2015). Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal*, *13*, 8–17. https://doi.org/10.1016/J.CSBJ.2014.11.005/ASSET/DF47FA96-4626-4074-BCF3-C254711D7688/MAIN.ASSETS/GR7.JPG

Lam, D., & Wunsch, D. C. (2014). *Clustering*. *1*, 1115–1149. https://doi.org/10.1016/B978-0-12-396502-8.00020-6

Langford, B. J., Daneman, N., Leung, V., & Langford, D. J. (2020). Cognitive bias: how understanding its impact on antibiotic prescribing decisions can help advance antimicrobial stewardship. *JAC-Antimicrobial Resistance*, *2*(4). https://doi.org/10.1093/JACAMR/DLAA107

Lawal, K., & Rafsanjani, H. N. (2022). Trends, benefits, risks, and challenges of IoT implementation in residential and commercial buildings. *Energy and Built Environment*, *3*(3), 251–266. https://doi.org/10.1016/J.ENBENV.2021.01.009

Le, H. M., Jiang, N., Agarwal, A., Dudík, M., Yue, Y., & Daumé, H. (2018). Hierarchical Imitation and Reinforcement Learning. *35th International Conference on Machine Learning, ICML 2018*, *7*, 4560–4573. https://arxiv.org/abs/1803.00590v2

Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature 2015 521:7553*, *521*(7553), 436–444. https://doi.org/10.1038/nature14539

Lee, S., & Karava, P. (2020). Towards smart buildings with self-tuned indoor thermal environments – A critical review. *Energy and Buildings*, *224*, 110172. https://doi.org/10.1016/J.ENBUILD.2020.110172

Lee, Y. T., Hsiao, W. H., Lin, Y. S., & Chou, S. C. T. (2017). Privacy-preserving data analytics in cloud-based smart home with community hierarchy. *IEEE Transactions on Consumer Electronics*, *63*(2), 200–207. https://doi.org/10.1109/TCE.2017.014777

Liang, J., Miao, H., Li, K., Tan, J., Wang, X., Luo, R., & Jiang, Y. (2025). A Review of Multi-Agent Reinforcement Learning Algorithms. *Electronics 2025, Vol. 14, Page 820*, *14*(4), 820. https://doi.org/10.3390/ELECTRONICS14040820

Lin, J., Ma, J., & Zhu, J. (2022). A Privacy-Preserving Federated Learning Method for Probabilistic Community-Level Behind-the-Meter Solar Generation Disaggregation. *IEEE Transactions on Smart Grid*, *13*(1). https://doi.org/10.1109/TSG.2021.3115904

Lowitzsch, J., Hoicka, C. E., & van Tulder, F. J. (2020). Renewable energy communities under the 2019 European Clean Energy Package – Governance model for the energy clusters of the future? *Renewable and Sustainable Energy Reviews*, *122*, 109489. https://doi.org/10.1016/J.RSER.2019.109489

Lu, S., Christie, G. A., Nguyen, T. T., Freeman, J. D., & Hsu, E. B. (2022). Applications of Artificial Intelligence and Machine Learning in Disasters and Public Health Emergencies. *Disaster Medicine and Public Health Preparedness*, *16*(4), 1674–1681. https://doi.org/10.1017/DMP.2021.125

Lyu, J., Wan, L., Li, X., & Lu, Z. (2024). Off-policy RL algorithms can be sample-efficient for continuous control via sample multiple reuse. *Information Sciences*, *666*, 120371. https://doi.org/10.1016/J.INS.2024.120371

Madrazo, F., Haas, T., Shen, Y.-S., Eugenio-Gozalbo, M., Landa Oregi, I., Urra-Uriarte, S., Gonzalez Ochoantesana, I., Anaya Rodríguez, M., & Molina-Costa, P. (2025). Enhancing Citizen Participation in Citizen-Centered Smart Cities: Insights from Two European Case Studies. *Urban Science 2025, Vol. 9, Page 140*, *9*(5), 140. https://doi.org/10.3390/URBANSCI9050140

Maduta, C., Melica, G., D'Agostino, D., & Bertoldi, P. (2022). Towards a decarbonised building stock by 2050: The meaning and the role of zero emission buildings (ZEBs) in Europe. *Energy Strategy Reviews*, *44*, 101009. https://doi.org/10.1016/J.ESR.2022.101009

Mallik, S., & Gangopadhyay, A. (2023). Proactive and reactive engagement of artificial intelligence methods for education: a review. *Frontiers in Artificial Intelligence*, *6*, 1151391. https://doi.org/10.3389/FRAI.2023.1151391/XML/NLM

Marcham, A. (2021). What Is Edge Computing? *Understanding Infrastructure Edge Computing*, 3–19. https://doi.org/10.1002/9781119763260.CH2

Maruyama, Y. (2022). Categorical Artificial Intelligence: The Integration of Symbolic and Statistical AI for Verifiable, Ethical, and Trustworthy AI. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *13154 LNAI*, 127–138. https://doi.org/10.1007/978-3-030-93758-4_14

Mary Reena, K. E., Theckethil Mathew, A., & Jacob, L. (2015). An Occupancy Based Cyber-Physical System Design for Intelligent Building Automation. *Mathematical Problems in Engineering*, *2015*(1), 132182. https://doi.org/10.1155/2015/132182

Marzband, M., Ghadimi, M., Sumper, A., & Domínguez-García, J. L. (2014). Experimental validation of a real-time energy management system using multi-period gravitational search algorithm for microgrids in islanded mode. *Applied Energy*, *128*, 164–174. https://doi.org/10.1016/J.APENERGY.2014.04.056

Masthoff, J. (2004). Group modeling: Selecting a sequence of television items to suit a group of viewers. *User Modelling and User-Adapted Interaction*, *14*(1), 37–85. https://doi.org/10.1023/B:USER.0000010138.79319.FD/METRICS

Meppelink, C. S., Smit, E. G., Fransen, M. L., & Diviani, N. (2019). "I was Right about Vaccination": Confirmation Bias and Health Literacy in Online Health Information Seeking. *Journal of Health Communication*, *24*(2), 129–140. https://doi.org/10.1080/10810730.2019.1583701

Merabet, G. H., Essaaidi, M., Haddou, M. Ben, Qolomany, B., Qadir, J., Anan, M., Al-Fuqaha, A., Abid, M. R., & Benhaddou, D. (2021). *Intelligent Building Control Systems for Thermal*

Comfort and Energy-Efficiency: A Systematic Review of Artificial Intelligence-Assisted Techniques.

Montesinos López, O. A., Montesinos López, A., & Crossa, J. (2022). Overfitting, Model Tuning, and Evaluation of Prediction Performance. *Multivariate Statistical Machine Learning Methods for Genomic Prediction*, 109–139. https://doi.org/10.1007/978-3-030-89010-0_4

Morimoto, J., & Ponton, F. (2021). Virtual reality in biology: could we become virtual naturalists? *Evolution: Education and Outreach*, *14*(1), 1–13. https://doi.org/10.1186/S12052-021-00147-X/FIGURES/4

Nabizadeh Rafsanjani, H., & Nabizadeh, A. H. (2023). Towards human-centered artificial intelligence (AI) in architecture, engineering, and construction (AEC) industry. *Computers in Human Behavior Reports*, *11*, 100319. https://doi.org/10.1016/J.CHBR.2023.100319

Nasri, E., Kania, R., & Tsauri, S. (2020). *Network Integration and Security Using IDS and Tunneling Methods*. 91–94. https://doi.org/10.2991/ASSEHR.K.200303.023

Neshenko, N., Bou-Harb, E., Crichigno, J., Kaddoum, G., & Ghani, N. (2019). Demystifying IoT Security: An Exhaustive Survey on IoT Vulnerabilities and a First Empirical Look on Internet-Scale IoT Exploitations. *IEEE Communications Surveys and Tutorials*, *21*(3), 2702–2733. https://doi.org/10.1109/COMST.2019.2910750

Nik, V. M., & Moazami, A. (2021). Using collective intelligence to enhance demand flexibility and climate resilience in urban areas. *Applied Energy*, *281*, 116106. https://doi.org/10.1016/J.APENERGY.2020.116106

Nishi, H., & Nakamura, Y. (2020). IoT-based monitoring for smart community. *Urban Systems Design: Creating Sustainable Smart Cities in the Internet of Things Era*, 335–344. https://doi.org/10.1016/B978-0-12-816055-8.00010-5

Noaman, M., Khan, M. S., Abrar, M. F., Ali, S., Alvi, A., & Saleem, M. A. (2022). Challenges in Integration of Heterogeneous Internet of Things. *Scientific Programming*, *2022*(1), 8626882. https://doi.org/10.1155/2022/8626882

Norval, C., & Singh, J. (2024). A Room with an Overview: Toward Meaningful Transparency for the Consumer Internet of Things. *IEEE Internet of Things Journal*, *11*(5), 7583–7603. https://doi.org/10.1109/JIOT.2023.3318369

Nozari, S., Krayani, A., Marin-Plaza, P., Marcenaro, L., Gomez, D. M., & Regazzoni, C. (2022). Active Inference Integrated With Imitation Learning for Autonomous Driving. *IEEE Access*, *10*, 49738–49756. https://doi.org/10.1109/ACCESS.2022.3172712

Olawade, D. B., Wada, O. Z., David-Olawade, A. C., Fapohunda, O., Ige, A. O., & Ling, J. (2024). Artificial intelligence potential for net zero sustainability: Current evidence and prospects. *Next Sustainability*, *4*, 100041. https://doi.org/10.1016/J.NXSUST.2024.100041

Oliveira, P., Novais, P., & Matos, P. (2018). Challenges in Smart Spaces: Aware of Users, Preferences, Behaviours and Habits. *Advances in Intelligent Systems and Computing*, *619*, 268–271. https://doi.org/10.1007/978-3-319-61578-3_34

Oliveira, T. A., Oliver, M., & Ramalhinho, H. (2020). Challenges for Connecting Citizens and Smart Cities: ICT, E-Governance and Blockchain. *Sustainability 2020, Vol. 12, Page 2926*, *12*(7), 2926. https://doi.org/10.3390/SU12072926

O'Malley, C., de Mars, P., Badesa, L., & Strbac, G. (2023). Reinforcement learning and mixed-integer programming for power plant scheduling in low carbon systems: Comparison and hybridisation. *Applied Energy*, *349*, 121659. https://doi.org/10.1016/J.APENERGY.2023.121659

Pardo-Cabello, A. J., Manzano-Gamero, V., & Puche-Cañas, E. (2022). Placebo: a brief updated review. *Naunyn-Schmiedeberg's Archives of Pharmacology 2022 395:11*, *395*(11), 1343–1356. https://doi.org/10.1007/S00210-022-02280-W

Park, S. Y., Kuo, P. Y., Barbarin, A., Kaziunas, E., Chow, A., Singh, K., Wilcox, L., & Lasecki, W. S. (2019). Identifying challenges and opportunities in human-AI collaboration in healthcare. *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*, 506–510. https://doi.org/10.1145/3311957.3359433

Pathmabandu, C., Grundy, J., Chhetri, M. B., & Baig, Z. (2023). Privacy for IoT: Informed consent management in Smart Buildings. *Future Generation Computer Systems*, *145*, 367–383. https://doi.org/10.1016/J.FUTURE.2023.03.045

Peiris, S., Lai, J. H. K., Kumaraswamy, M. M., & Hou, H. (Cynthia). (2023). Smart retrofitting for existing buildings: State of the art and future research directions. *Journal of Building Engineering*, *76*, 107354. https://doi.org/10.1016/J.JOBE.2023.107354

Perera, A. T. D., & Kamalaruban, P. (2021). Applications of reinforcement learning in energy systems. *Renewable and Sustainable Energy Reviews*, *137*, 110618. https://doi.org/10.1016/J.RSER.2020.110618

Pigozzi, G., Tsoukiàs, A., & Viappiani, P. (2015). Preferences in artificial intelligence. *Annals of Mathematics and Artificial Intelligence 2015 77:3*, *77*(3), 361–401. https://doi.org/10.1007/S10472-015-9475-5

Poo, D., Chng, B., & Goh, J. M. (2003). A hybrid approach for user profiling. *Proceedings of the 36th Annual Hawaii International Conference on System Sciences, HICSS 2003*. https://doi.org/10.1109/HICSS.2003.1174242

Prandi, C., Monti, L., Ceccarini, C., & Salomoni, P. (2020). Smart Campus: Fostering the Community Awareness Through an Intelligent Environment. *Mobile Networks and Applications*, *25*(3), 945–952. https://doi.org/10.1007/S11036-019-01238-2/FIGURES/4

Qiu, T., Chen, N., Li, K., Atiquzzaman, M., & Zhao, W. (2018). How can heterogeneous internet of things build our future: A survey. *IEEE Communications Surveys and Tutorials*, *20*(3), 2011–2027. https://doi.org/10.1109/COMST.2018.2803740

Raman, R., Kowalski, R., Achuthan, K., Iyer, A., & Nedungadi, P. (2025). Navigating artificial general intelligence development: societal, technological, ethical, and brain-inspired pathways. *Scientific Reports 2025 15:1*, *15*(1), 1–22. https://doi.org/10.1038/s41598-025-92190-7

Ribeiro, B., Silva, R., Mota, B., Gomes, L., & Vale, Z. (2024). Learning-Based Models for Intelligent Control Over Air Conditioning Units in a Smart Building. *Lecture Notes in Networks and Systems*, *888 LNNS*, 197–207. https://doi.org/10.1007/978-3-031-75013-7_19

Rubí, J. N. S., & de Lira Gondim, P. R. (2021a). IoT-based platform for environment data sharing in smart cities. *International Journal of Communication Systems*, *34*(2), e4515. https://doi.org/10.1002/DAC.4515

Rubí, J. N. S., & de Lira Gondim, P. R. (2021b). IoT-based platform for environment data sharing in smart cities. *International Journal of Communication Systems*, *34*(2), e4515. https://doi.org/10.1002/DAC.4515

Schiller, E., Aidoo, A., Fuhrer, J., Stahl, J., Ziörjen, M., & Stiller, B. (2022). Landscape of IoT security. *Computer Science Review*, *44*, 100467. https://doi.org/10.1016/J.COSREV.2022.100467

Schmidt, G., Lappus, G., Freyberger, F., & Dungern, O. V. (1987). Communication and Networking in Control Education. *IFAC Proceedings Volumes*, *20*(5), 247–252. https://doi.org/10.1016/S1474-6670(17)55272-6

Schürmann, T., & Beckerle, P. (2020). Personalizing Human-Agent Interaction Through Cognitive Models. *Frontiers in Psychology*, *11*, 561510. https://doi.org/10.3389/FPSYG.2020.561510/BIBTEX

Shalev-Shwartz, S., & Ben-David, S. (2014). Understanding Machine Learning: From Theory to Algorithms. *Understanding Machine Learning: From Theory to Algorithms*, *9781107057135*, 1–397. https://doi.org/10.1017/CBO9781107298019

Shin, C., Dey, A. K., & Woo, W. (2010). TOWARD COMBINING AUTOMATIC RESOLUTION WITH SOCIAL MEDIATION FOR RESOLVING MULTIUSER CONFLICTS. *Cybernetics and Systems: An International Journal*, *41*(2), 146–166. https://doi.org/10.1080/01969720903584282

Shrestha, S., Mount, J., Vald, G., Sermet, Y., Samuel, D. J., Bryant, C., Peralta, A. C., Beck, M. W., Meyers, S. D., Muller-Karger, F. E., Cwiertny, D., & Demir, I. (2025). A community-centric intelligent cyberinfrastructure for addressing nitrogen pollution using web systems and conversational AI. *Environmental Science & Policy*, *167*, 104055. https://doi.org/10.1016/J.ENVSCI.2025.104055

Singh, K., Kaushik, K., Ahatsham, & Shahare, V. (2020). Role and Impact of Wearables in IoT Healthcare. *Advances in Intelligent Systems and Computing*, *1090*, 735–742. https://doi.org/10.1007/978-981-15-1480-7_67

Spillo, G., Musto, C., de Gemmis, M., Lops, P., & Semeraro, G. (2024). Recommender systems based on neuro-symbolic knowledge graph embeddings encoding first-order logic rules. *User Modeling and User-Adapted Interaction*, *34*(5), 2039–2083. https://doi.org/10.1007/S11257-024-09417-X/TABLES/12

Sugandha, Freestone, R., & Favaro, P. (2022). The social sustainability of smart cities: A conceptual framework. *City, Culture and Society*, *29*, 100460. https://doi.org/10.1016/J.CCS.2022.100460

Sun, Y., Peng, M., Zhou, Y., Huang, Y., & Mao, S. (2019). Application of Machine Learning in Wireless Networks: Key Techniques and Open Issues. *IEEE Communications Surveys and Tutorials*, *21*(4), 302–3108. https://doi.org/10.1109/COMST.2019.2924243

Taghikhah, F., Erfani, E., Bakhshayeshi, I., Tayari, S., Karatopouzis, A., & Hanna, B. (2022). Artificial intelligence and sustainability: solutions to social and environmental challenges. *Artificial Intelligence and Data Science in Environmental Sensing*, 93–108. https://doi.org/10.1016/B978-0-323-90508-4.00006-X

Terenzi, D., Liu, L., Bellucci, G., & Park, S. Q. (2021). Determinants and modulators of human social decisions. *Neuroscience & Biobehavioral Reviews*, *128*, 383–393. https://doi.org/10.1016/J.NEUBIOREV.2021.06.041

To, W. M., Lai, L. S. L., Lam, K. H., & Chung, A. W. L. (2018). Perceived Importance of Smart and Sustainable Building Features from the Users' Perspective. *Smart Cities 2018, Vol. 1, Pages 163-175*, *1*(1), 163–175. https://doi.org/10.3390/SMARTCITIES1010010

Tomor, Z., Meijer, A., Michels, A., & Geertman, S. (2019). Smart Governance For Sustainable Cities: Findings from a Systematic Literature Review. *Journal of Urban Technology*, *26*(4), 3–27. https://doi.org/10.1080/10630732.2019.1651178

Tushar, W., Saha, T. K., Yuen, C., Morstyn, T., McCulloch, M. D., Poor, H. V., & Wood, K. L. (2019). A motivational game-theoretic approach for peer-to-peer energy trading in the smart grid. *Applied Energy*, *243*, 10–20. https://doi.org/10.1016/J.APENERGY.2019.03.111

Tushar, W., Yuen, C., Saha, T. K., Morstyn, T., Chapman, A. C., Alam, M. J. E., Hanif, S., & Poor, H. V. (2021). Peer-to-peer energy systems for connected communities: A review of recent advances and emerging challenges. *Applied Energy*, *282*, 116131. https://doi.org/10.1016/J.APENERGY.2020.116131

Uslu, S., Kaur, D., Rivera, S. J., Durresi, A., Babbar-Sebens, M., & Tilt, J. H. (2021). A Trustworthy Human–Machine framework for collective decision making in Food–Energy–Water management: The role of trust sensitivity. *Knowledge-Based Systems*, *213*, 106683. https://doi.org/10.1016/J.KNOSYS.2020.106683 van Delden, S. H., SharathKumar, M., Butturini, M., Graamans, L. J. A., Heuvelink, E., Kacira, M., Kaiser, E., Klamer, R. S., Klerkx, L., Kootstra, G., Loeber, A., Schouten, R. E., Stanghellini, C., van Ieperen, W., Verdonk, J. C., Vialet-Chabrand, S., Woltering, E. J., van de Zedde, R., Zhang, Y., & Marcelis, L. F. M. (2021). Current status and future challenges in implementing and upscaling vertical farming systems. *Nature Food 2021 2:12*, *2*(12), 944–956. https://doi.org/10.1038/s43016-021-00402-w

Verma, A., Prakash, S., Srivastava, V., Kumar, A., & Mukhopadhyay, S. C. (2019). Sensing, Controlling, and IoT Infrastructure in Smart Building: A Review. *IEEE Sensors Journal*, *19*(20), 9036–9046. https://doi.org/10.1109/JSEN.2019.2922409

Von Frankenberg, N. (2021). Towards Resolving Thermal Comfort Conflicts in Shared Spaces. *Conference on Human Factors in Computing Systems - Proceedings*. https://doi.org/10.1145/3411763.3451512

Wang, F. Y., Zeng, D., Carley, K. M., & Mao, W. (2007). Social computing: From social informatics to social intelligence. *IEEE Intelligent Systems*, *22*(2), 79–83. https://doi.org/10.1109/MIS.2007.41

Wang, R., Member, S., Li, H., Member, S., & Liu, E. (2021). *Blockchain-Based Federated Learning in Mobile Edge Networks with Application in Internet of Vehicles*. 1. https://arxiv.org/abs/2103.01116v1

Wang, R., Wu, Z., Lou, J., & Jiang, Y. (2022). Attention-based dynamic user modeling and Deep Collaborative filtering recommendation. *Expert Systems with Applications*, *188*, 116036. https://doi.org/10.1016/J.ESWA.2021.116036

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning 1992 8:3*, *8*(3), 279–292. https://doi.org/10.1007/BF00992698

Weiser, M. (1991). The Computer for the 21 st Century. *Scientific American*. https://www.jstor.org/stable/24938718

Xia, F., Liu, L., Li, J., Ma, J., & Vasilakos, A. V. (2015). Socially aware networking: A survey. *IEEE Systems Journal*, *9*(3), 904–921. https://doi.org/10.1109/JSYST.2013.2281262

Xiao, Y., Shi, G., Li, Y., Saad, W., & Poor, H. V. (2020). Toward Self-Learning Edge Intelligence in 6G. *IEEE Communications Magazine*, *58*(12), 34–40. https://doi.org/10.1109/MCOM.001.2000388

Xiong, L., Li, P., Wang, Z., & Wang, J. (2019). *Multi-agent based multi objective renewable energy management for diversified community power consumers*. https://doi.org/10.1016/j.apenergy.2019.114140

Yang, D., Li, D., Chen, Y., Siano, P., Lim, S. B., & Yigitcanlar, T. (2022). Participatory Governance of Smart Cities: Insights from e-Participation of Putrajaya and Petaling Jaya, Malaysia.

Smart Cities 2022, Vol. 5, Pages 71-89, *5*(1), 71–89. https://doi.org/10.3390/SMARTCITIES5010005

Yang, Q., & Wang, H. (2021). Distributed energy trading management for renewable prosumers with HVAC and energy storage. *Energy Reports*, *7*, 2512–2525. https://doi.org/10.1016/J.EGYR.2021.03.038

Yang, Z., Ghadamyari, M., Khorramdel, H., Seyed Alizadeh, S. M., Pirouzi, S., Milani, M., Banihashemi, F., & Ghadimi, N. (2021). Robust multi-objective optimal design of islanded hybrid system with renewable and diesel sources/stationary and mobile energy storage systems. *Renewable and Sustainable Energy Reviews*, *148*, 111295. https://doi.org/10.1016/J.RSER.2021.111295

Yau, K. L. A., Qadir, J., Khoo, H. L., Ling, M. H., & Komisarczuk, P. (2017). A Survey on Reinforcement Learning Models and Algorithms for Traffic Signal Control. *ACM Computing Surveys (CSUR)*, *50*(3). https://doi.org/10.1145/3068287

Ying, Q. F., Chiu, D. M., Venkatramanan, S., & Zhang, X. (2018). User modeling and usage profiling based on temporal posting behavior in OSNs. *Online Social Networks and Media*, *8*, 32–41. https://doi.org/10.1016/J.OSNEM.2018.10.003

Ylonen, T., Turner, P., Scarfone, K., & Souppaya, M. (n.d.). *Security of Interactive and Automated Access Management Using Secure Shell (SSH)*. https://doi.org/10.6028/NIST.IR.7966

Yorgos, Y. J. S., Golias, M., Dedes, G., Douligeris, C., & Mishra, S. (2019). Challenges, Risks and Opportunities for Connected Vehicle Services in Smart Cities and Communities. *IFAC-PapersOnLine*, *51*(34), 139–144. https://doi.org/10.1016/J.IFACOL.2019.01.056

Zador, A., Escola, S., Richards, B., Ölveczky, B., Bengio, Y., Boahen, K., Botvinick, M., Chklovskii, D., Churchland, A., Clopath, C., DiCarlo, J., Ganguli, S., Hawkins, J., Körding, K., Koulakov, A., LeCun, Y., Lillicrap, T., Marblestone, A., Olshausen, B., … Tsao, D. (2023). Catalyzing next-generation Artificial Intelligence through NeuroAI. *Nature Communications 2023 14:1*, *14*(1), 1–7. https://doi.org/10.1038/s41467-023-37180-x

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, *23*, 100224. https://doi.org/10.1016/J.JII.2021.100224

Zhang, K. (2021). Research on Building Intelligent Technology Based on Modern Information Artificial Intelligence Building Concept. *Journal of Physics: Conference Series*, *2037*(1), 012101. https://doi.org/10.1088/1742-6596/2037/1/012101

Zhang, Y., Chen, Y., Wang, J., & Pan, Z. (2023). Unsupervised Deep Anomaly Detection for Multi-Sensor Time-Series Signals. *IEEE Transactions on Knowledge and Data Engineering*, *35*(2), 2118–2132. https://doi.org/10.1109/TKDE.2021.3102110

Zhang, Y., Qu, Y., Gao, L., Luan, T. H., Jolfaei, A., & Zheng, J. X. (2023). Privacy-preserving data analytics for smart decision-making energy systems in sustainable smart community. *Sustainable Energy Technologies and Assessments*, *57*, 103144. https://doi.org/10.1016/J.SETA.2023.103144

Zheng, H., Paiva, A. R., & Gurciullo, C. S. (2020). *Advancing from Predictive Maintenance to Intelligent Maintenance with AI and IIoT*. https://arxiv.org/abs/2009.00351v1

Zhu, M., Pan, Y., Wu, Z., Xie, J., Huang, Z., & Kosonen, R. (2021). An occupant-centric air-conditioning system for occupant thermal preference recognition control in personal micro-environment. *Building and Environment*, *196*, 107749. https://doi.org/10.1016/J.BUILDENV.2021.107749

Zieglmeier, V., & Pretschner, A. (2021). *Trustworthy Transparency by Design*. https://arxiv.org/abs/2103.10769v2

Zieni, R., Massari, L., & Calzarossa, M. C. (2023). Phishing or Not Phishing? A Survey on the Detection of Phishing Websites. *IEEE Access*, *11*, 18499–18519. https://doi.org/10.1109/ACCESS.2023.3247135