COMPUTATIONAL AI PREDICTIVE MODELS FOR SUSTAINABLE ENERGY AND WASTE MANAGEMENT

By Mattew Aderinsola Olawumi

A THESIS

Presented to the Department of
Computational Intelligence
program at Selinus University

Faculty of Computer Science
in fulfillment of the requirements
for the degree of Doctor of Philosophy
In Computational Intelligence

2025

Author's declaration

I, **Mattew Aderinsola Olawumi**, hereby declare that this thesis, submitted in partial fulfilment of the requirements for the award of the degree of Doctor of Philosophy, is the result of my own original work. I confirm that I am the sole and corresponding author of the research presented, and that it has not been submitted, in whole or in part, for the award of any other degree or qualification at this or any other institution.

The conception, design, execution, and interpretation of the study were carried out under my direct supervision and responsibility. While co-authors provided supportive contributions in related publications, they did not undertake substantial technical or intellectual input into the work presented in this thesis.

All sources of data, information, and scholarly ideas drawn from the work of others have been clearly acknowledged and fully referenced in accordance with academic standards.

Acknowledgement

I would like to express my deepest gratitude to all those who supported and contributed to the successful completion of this doctoral research.

First and foremost, I am profoundly grateful to **Professor Salvatore Fava**, whose academic guidance, insightful feedback, and unwavering support provided the intellectual foundation for this thesis. His expertise and commitment to excellence continually challenged me to refine my thinking and pursue research of the highest standard.

I extend my sincere appreciation to **Dr. Bankole I. Oladapo** for his invaluable mentorship, supervision, and encouragement throughout this journey. His practical advice, scholarly direction, and consistent support played a vital role in shaping the methodology, analysis, and overall coherence of this work.

Special thanks are due to the wider **university team** for their administrative assistance, access to research facilities, and for fostering an enabling environment that made this work possible.

I am also indebted to the **participants** who generously shared their time, experiences, and perspectives during the surveys and interviews. Their openness and willingness to contribute formed the heart of this research, without which this thesis could not have been realised.

On a personal note, I wish to acknowledge my beloved **wife, Temitope**, and my **daughter, Tomisin**, whose constant belief in me inspired resilience and perseverance. Their love and encouragement have been an unfailing source of motivation.

To my **parents, Mr. and Mrs. Olawumi**, I owe a special debt of gratitude for their unwavering support, sacrifices, and encouragement, which have been the foundation of my academic journey.

Finally, I am thankful to my **friends and colleagues** for their emotional support, patience, and encouragement during the most challenging phases of this work. Their belief in me strengthened my resolve and carried me through to completion.

To all of you, I offer my heartfelt thanks.

Dedication

This work is lovingly dedicated to my beloved **wife, Temitope**, and my cherished **daughter, Tomisin**—the true foundation of my academic journey.

Your unwavering belief in me, your patience, and your words of encouragement sustained me through the most challenging moments. Your love, strength, and faith have been my anchor, inspiring resilience and giving me the courage to persevere.

This achievement is as much yours as it is mine.

Table of Contents

List of Figures

List of Table

Abstract

Global sustainability challenges, particularly those posed by climate change, resource scarcity, and environmental degradation, necessitate innovative solutions that transcend traditional approaches. Energy consumption and waste generation remain two of the most critical contributors to carbon emissions and ecological disruption, yet they are often studied and managed in isolation. This PhD research addresses this gap by developing and validating an integrated computational artificial intelligence (AI) framework that simultaneously models energy demand and waste generation to support Net Zero and circular economy transitions.

The study applies a suite of machine learning and deep learning techniques—including multilayer perceptrons (MLP), long short-term memory (LSTM) networks, Gradient Boosting, and Random Forests—to large, heterogeneous datasets derived from smart meters, municipal waste records, IoT-enabled environmental sensors, and weather data. These models capture nonlinear dependencies and temporal dynamics that traditional rule-based or statistical approaches, such as ARIMA, often fail to recognise. Comparative evaluations demonstrate that AI-driven predictive systems substantially improve forecasting accuracy and operational responsiveness across both energy and waste domains.

A distinctive contribution of this research lies in the integration of explainable AI (XAI) methods, including SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations). These tools enhance transparency and accountability by elucidating how input features influence predictions, thereby fostering trust among policymakers, utility providers, and municipal authorities. This commitment to interpretability is particularly significant given the high stakes of sustainability decision-making and the ethical considerations surrounding data use, privacy, and algorithmic bias.

The methodological framework is grounded in pragmatic positivism, combining rigorous quantitative model development with real-world case studies and simulations. Findings confirm that AI-enhanced predictive systems can optimise energy demand response strategies, improve waste collection scheduling, and facilitate resource-efficient planning. Moreover, by linking energy and waste dynamics within a single architecture, the framework highlights synergies such as waste-to-energy potential and cross-sectoral policy impacts.

This thesis makes three major contributions: (1) advancing predictive modelling techniques for energy and waste systems; (2) embedding explainability and ethical considerations within AI frameworks to enhance adoption and stakeholder confidence; and (3) proposing a transferable, scalable decision-support platform aligned with the United Nations Sustainable Development Goals, particularly SDG 7 (Affordable and Clean Energy), SDG 12 (Responsible Consumption and Production), and SDG 13 (Climate Action).

Ultimately, the research demonstrates that AI-driven predictive frameworks are not merely technological enhancements but transformative enablers of sustainability, capable of guiding societies towards resilient, low-carbon, and circular futures.

Keywords: Artificial Intelligence (AI); Machine Learning (ML); Deep Learning (LSTM, MLP); Explainable AI (XAI); Energy Demand Forecasting; Waste Management Optimisation; Circular Economy / Net Zero

Chapter 1

1 INTRODUCTION

1.1 Background and Context

The world today stands at a critical juncture, where the pursuit of sustainable development has become not merely desirable, but essential [1,2]. The dual challenges of climate change and environmental degradation have brought into sharp focus the urgent need for innovative solutions that can balance economic growth with environmental stewardship. Nowhere is this tension more visible than in the domains of energy consumption and waste management. Both sectors, although vital to modern life, have historically been among the most significant contributors to carbon emissions, resource depletion, and ecological disruption [3,4]. Over the last few decades, global energy demand has continued to surge, driven by industrialisation, urbanisation, and rising living standards, particularly in developing economies. According to reports from bodies such as the International Energy Agency (IEA), energy consumption in residential, commercial, and industrial sectors has reached unprecedented levels, placing enormous pressure on national grids and accelerating the depletion of non-renewable energy sources. This has been paralleled by increasing volumes of waste generation, exacerbated by consumerism, population growth, and inadequate recycling infrastructure in many regions [5,6].

Governments, industries, and communities have responded in various ways, promoting the adoption of renewable energy, mandating recycling targets, and investing in infrastructure designed to mitigate environmental harm [7,8]. Yet, despite these efforts, significant gaps remain. Energy grids continue to experience instability, particularly as intermittent renewable energy sources, such as wind and solar, become more prevalent. Similarly, waste management systems often struggle to cope with fluctuating recycling rates, contamination in waste streams, and inefficient collection schedules, resulting in missed targets and an increased reliance on landfills [9,10]. It is within this complex and evolving context that the application of artificial intelligence (AI) and computational intelligence techniques offers a promising avenue. AI, particularly machine learning (ML) and its advanced forms such as deep learning, has demonstrated remarkable success in fields ranging from healthcare diagnostics to financial forecasting. What makes AI particularly compelling in the context of energy and waste management is its ability to model complex, non-linear systems, learn from historical and real-time data, and make predictions or decisions that are adaptive rather than static [11, 12].

Historically, energy management systems have relied heavily on rule-based approaches. These systems operate on pre-defined thresholds, schedules, or heuristics, often without accounting for the dynamic interplay between variables such as weather conditions, occupancy patterns, or fluctuating energy prices. Similarly, traditional waste management planning has often involved fixed collection routes or schedules that do not account for actual waste generation rates or changing recycling behaviours across different neighbourhoods or periods [13,14]. Such approaches, although operationally simple, fall short in terms of efficiency and responsiveness. They fail to capitalise on the rich data streams now available from smart meters, sensors, Internet of Things (IoT) devices, and other modern technologies. Here, AI steps in with the potential to transform these systems. By processing large volumes of diverse data and learning complex patterns, AI models can provide far more accurate forecasts, optimise operations in real-time, and adapt to changing conditions in ways that conventional systems cannot [15, 16].

This PhD research emerges at the intersection of these critical global challenges and technological possibilities. The work is grounded in the belief that computational intelligence can deliver robust, scalable, and explainable solutions that address the limitations of current energy and waste systems. It explores the design and deployment of predictive models that not only offer higher accuracy but also deliver insights into the decision-making process through the use of explainable AI (XAI) techniques. In a world where black-box models can hinder trust and adoption, transparency is crucial, particularly in public-sector applications where accountability is paramount.

The growing momentum towards Net Zero targets further shapes the context for this research. Nations worldwide, from the United Kingdom to emerging economies, have set ambitious goals to achieve carbon neutrality by mid-century. Meeting these targets demands a rethinking of how energy is generated, distributed, and consumed, as well as how materials are recovered, and waste is minimised. Incremental improvements alone will not suffice; systemic change is required. This thesis argues that AI-driven predictive modelling can catalyse such change [17,18]. In terms of energy management, this work focuses particularly on the residential and small commercial building sectors, which collectively account for a significant share of global energy use. It investigates the use of AI models, including multilayer perceptrons (MLPs) and long short-term memory (LSTM) networks, to forecast energy demand and optimise demand response strategies. A critical contribution lies in comparing these advanced models to traditional rule-based methods, quantifying the performance gains, and identifying the trade-offs in terms of complexity, computational requirements, and interpretability [19, 20].

The waste management component of the research similarly addresses key inefficiencies in existing systems. By applying LSTM models to recycling rate data, this thesis demonstrates how AI can provide more reliable forecasts, enabling municipalities to plan more effective collection schedules and resource allocations. The ability of such models to capture temporal dependencies, showing how patterns from previous weeks or months may influence today's recycling behaviour, offers a distinct advantage over conventional statistical techniques [21,22].

Another critical strand running through this research is the integration of explainable AI. As predictive models become more complex, there is a risk that they turn into opaque black boxes, undermining user trust and limiting their practical application. The thesis tackles this challenge by incorporating techniques such as SHAP (Shapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations), which help to clarify how specific inputs influence the model's predictions. This not only promotes transparency but also supports better decision-making by human operators, who remain accountable for their outcomes.

At a broader level, this research is situated within the global shift towards smarter, data-driven cities. As urban areas grow and digitise, there is a pressing need for integrated solutions that can simultaneously address multiple sustainability challenges. The convergence of AI with IoT and smart grid technologies represents a powerful enabler in this regard. The thesis contributes to this vision by demonstrating how AI models can integrate data from diverse sources, such as smart meters, weather stations, and waste collection records, and transform it into actionable intelligence that benefits both providers and end-users [23, 24].

Furthermore, the work recognises the importance of scalability and adaptability. While much of the empirical work is grounded in case studies or simulations focused on particular regions or building types, the modelling frameworks are designed with generalisability in mind. The methods developed could, with appropriate local data, be applied to other settings, whether in dense urban centres or more rural communities. This flexibility is key, given the varied contexts in which energy and waste challenges manifest [25,26]. It is also worth noting that this research does not view AI as a panacea. Instead, it acknowledges the limitations and risks associated with algorithmic decision-making. Issues such as data privacy, cybersecurity, and potential biases in training data are treated not as peripheral concerns but as central considerations in model development and evaluation. The thesis advocates for an ethical and responsible approach to AI deployment, one that aligns technological innovation with societal values and the public good [27, 28].

In summary, the background and context for this PhD research are shaped by a confluence of environmental urgency, technological opportunity, and policy imperatives. The work addresses real-world challenges with rigorously developed, computationally intelligent solutions that advance both academic knowledge and practical outcomes. It situates itself within a global movement towards sustainable, resilient, and transparent systems, positioning AI not as a futuristic ideal but as a present-day tool with transformative potential.

1.2 The Global Challenge of Sustainable Energy and Waste Management

Across the world, the intertwined issues of energy consumption and waste generation represent two of the most urgent and complex sustainability challenges of our time. At the heart of these challenges lies the pressing need to reconcile economic development with environmental protection. In both developed and developing regions, governments, industries, and communities are grappling with the consequences of unsustainable energy practices and ineffective waste management systems. These problems not only threaten ecological integrity but also undermine public health, economic stability, and social well-being [29, 30].

The energy demand continues to rise at an unprecedented rate, fuelled by population growth, urbanisation, technological advancement, and changing lifestyles. Global energy consumption has more than doubled since the 1970s, and projections indicate that this trend will persist in the coming decades, particularly as emerging economies expand their industrial bases and improve living standards. At the same time, the way we generate and use energy remains heavily reliant on fossil fuels. Despite the remarkable growth of renewable energy sources, coal, oil, and natural gas continue to account for the lion's share of the global energy supply. This dependence has led to soaring greenhouse gas emissions, which significantly contribute to climate change and its associated impacts, including rising sea levels and more frequent and severe weather events [31, 32]. Compounding this energy dilemma is the challenge of waste. The modern world generates waste on a scale unprecedented in human history. Municipal solid waste, industrial by-products, electronic waste, and agricultural residues all contribute to a global waste crisis. The World Bank estimates that by 2050, annual waste generation could rise to 3.4 billion tonnes if no meaningful interventions are made. Much of this waste ends up in landfills or, worse still, is openly dumped, polluting the land, rivers, and oceans. Plastic waste, in particular, has garnered widespread attention due to its persistence in the environment and harmful effects on marine ecosystems. Both energy and waste sectors are critical levers in the global effort to achieve sustainable development and mitigate climate change. The United Nations' Sustainable Development Goals (SDGs), particularly SDG 7 (Affordable and Clean Energy) and SDG 12 (Responsible Consumption and Production), emphasise the need to transition towards cleaner energy systems and more sustainable waste management practices [33,34].

Progress remains uneven and, in many cases, inadequate. The transition is not simply a matter of technological substitution or policy reform. It requires systemic change, a fundamental rethinking of how we produce, distribute, consume, and manage resources. One of the most significant barriers to this transition is the complexity and scale of the problem. Energy systems are vast and interconnected, involving multiple actors, infrastructures, and technologies [35, 36]. The variability of renewable energy sources, such as solar and wind power, introduces new challenges for maintaining grid stability and reliability. Waste management, meanwhile, operates at the interface of economic, social, and environmental systems, requiring coordination across local authorities, private firms, and the public. Both sectors face challenges related to financing, governance, and social acceptance that complicate efforts to implement sustainable solutions [37, 38].

The growing role of artificial intelligence and computational intelligence offers a new frontier in addressing these challenges. AI's capacity to process large volumes of data, identify patterns, and optimise complex systems aligns well with the needs of both energy and waste management sectors. In energy systems, predictive models can forecast demand with greater accuracy, enabling more effective balancing of supply and demand, and facilitating the integration of renewable sources. Smart grids equipped with AI can adapt in real-time to changes in consumption patterns, improving efficiency and resilience. Similarly, in waste management, AI can enhance planning and operations. Predictive analytics can forecast waste generation rates, allowing cities to design more responsive collection schedules. Machine learning models can support the sorting and classification of waste, improving recycling rates and reducing contamination. Moreover, AI-driven decision-support systems can help policymakers evaluate trade-offs between various waste treatment options, including recycling, composting, and energy recovery [39, 40].

However, the promise of AI must be seen in the context of broader systemic issues. Technology alone cannot resolve the global challenge of sustainable energy and waste management. There are significant concerns about data availability, quality, and privacy, especially in regions with limited infrastructure. AI models require robust, representative datasets to function effectively. Yet, in many parts of the world, such data are sparse or unreliable. Furthermore, the deployment of AI raises ethical questions about accountability, transparency, and potential biases in decision-making [41, 42].

Energy inequality is another dimension of the global challenge. While some countries invest in advanced renewable technologies and smart infrastructure, others struggle to provide even basic access to electricity. According to the International Energy Agency, nearly 770 million people worldwide still lack access to electricity, with the majority residing in sub-Saharan Africa. The push for cleaner, smarter energy systems must therefore also address the imperative of energy justice, ensuring that transitions do not deepen existing disparities. Waste management systems also reveal stark inequalities. In wealthier nations, waste collection and treatment are often taken for granted, supported by established infrastructures and regulatory frameworks [43, 44].

In contrast, low-income countries frequently rely on informal waste pickers, open dumping, and rudimentary landfills. The health and livelihoods of millions depend on these informal systems, which are often overlooked in formal policy discussions. A truly sustainable approach to waste management must recognise and integrate these realities, supporting inclusive strategies that improve outcomes for all stakeholders. Policy and governance frameworks play a decisive role in shaping the trajectory of energy and waste transitions. Ambitious national and international targets, such as commitments to Net-Zero emissions, provide essential signals, but translating these into action requires coherent strategies, institutional capacity, and political will. Market mechanisms, subsidies, regulations, and public investments all influence the pace and direction of change. Importantly, citizen engagement is critical. Public attitudes towards renewable energy installations, recycling practices, and waste infrastructure can either enable or hinder progress [45, 46]. The global COVID-19 pandemic has added another layer of complexity to these challenges. Lockdowns and economic disruptions led to shifts in energy consumption patterns, characterised by reduced industrial demand and increased residential use. The pandemic also generated new waste streams, such as personal protective equipment (PPE), highlighting gaps in existing waste systems. At the same time, recovery plans offer an opportunity to align economic stimulus with sustainability goals, supporting green energy investments and circular economy initiatives.

This PhD research, situated at the intersection of these global challenges and technological frontiers, seeks to contribute practical, data-driven solutions that advance both academic understanding and real-world impact. It examines how computational artificial intelligence can be harnessed to enhance the predictive capabilities of energy and waste systems, enabling more sustainable and efficient operations. The focus is not only on improving technical performance, for example, through more accurate forecasts or optimised schedules, but also on ensuring that the models developed are transparent, explainable, and aligned with societal values. A particular strength of this research lies in its integrative approach. It recognises that energy and waste challenges cannot be addressed in isolation. The thesis examines how AI models can leverage diverse data sources, including smart meters, weather sensors, waste collection records, and others, to develop comprehensive solutions. For instance, by linking energy consumption forecasts with occupancy patterns and weather conditions, predictive models can enable more innovative demand response strategies that reduce strain on the grid while maintaining user comfort. Similarly, by analysing temporal trends in recycling rates, AI can support dynamic planning that adjusts to changing behaviours and seasons [47, 48].

The work also acknowledges the need for flexibility and scalability. The proposed solutions are designed to be adaptable to different contexts, ranging from dense urban centres with advanced infrastructures to smaller towns or regions where resources may be more limited. This reflects an understanding that sustainability is not a one-size-fits-all proposition; local conditions, cultures, and capacities matter. The aim is to provide frameworks and tools that can be customised to suit varying needs and priorities. In addition to its technical contributions, the research engages with the broader questions of governance, ethics, and inclusion. It addresses how AI can support, rather than replace, human decision-making, offering insights and recommendations that enhance the capabilities of planners, engineers, and policymakers. It explores how explainable AI techniques can build trust among stakeholders, ensuring that the adoption of advanced technologies does not come at the cost of transparency or accountability. It also highlights the importance of considering the social dimensions of technological change, from data privacy to the impacts on jobs and livelihoods. Ultimately, the global challenge of sustainable energy and waste management is a defining issue of our era. It tests our collective ability to innovate, collaborate, and act with foresight [49, 50]. The scale and urgency of the problem demand ambitious solutions that harness the best of human ingenuity and technological possibility. This thesis contributes to that effort, offering evidence-based, computationally intelligent models that can help chart a more sustainable and equitable future. Figure 1 illustrates the integration of AI-driven predictive models with sustainable energy and waste management systems.

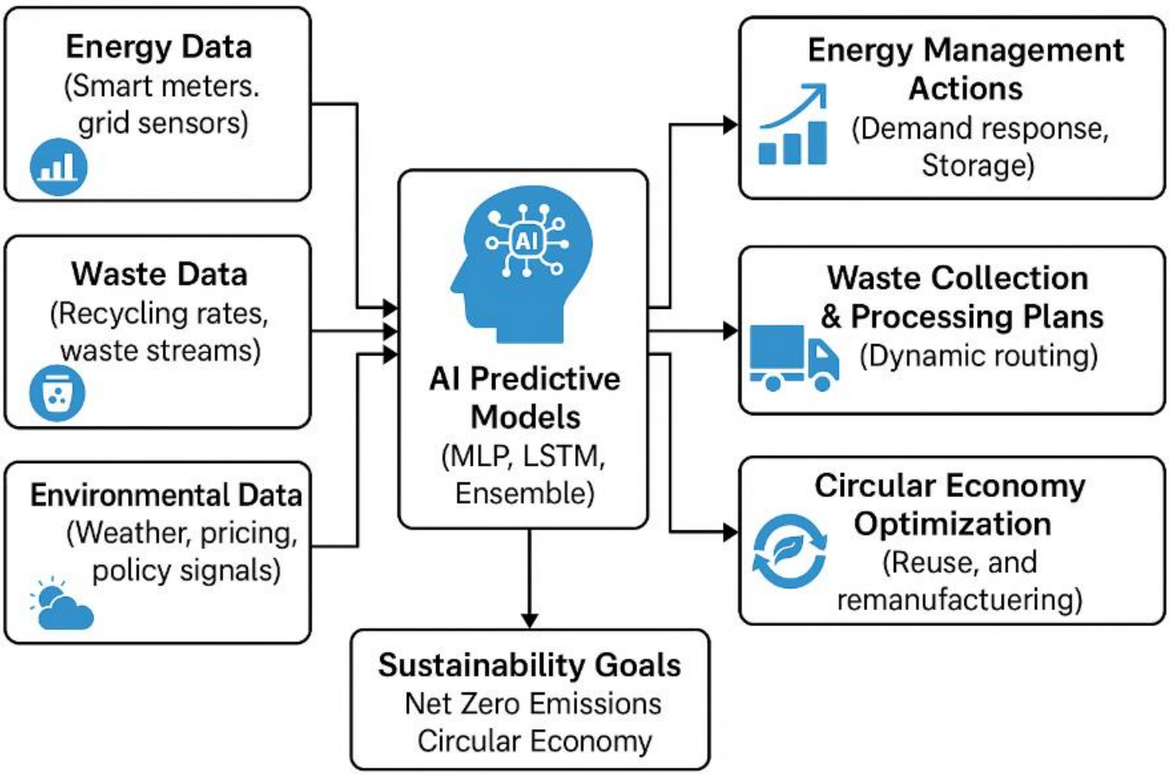

Figure1.1 Conceptual Framework of the Research.

1.3 The Role of Artificial Intelligence in Achieving Net Zero and Circular Economy Goals

The transition towards a sustainable, low-carbon future is perhaps the most ambitious and complex undertaking of our time. Governments, businesses, and communities worldwide are grappling with the dual imperatives of drastically reducing carbon emissions while reshaping economic systems to be more resource-efficient, resilient, and equitable. Central to this vision are two interlinked concepts: achieving *Net Zero* emissions and embracing the *circular economy*. The former focuses on balancing the amount of greenhouse gases emitted with those removed from the atmosphere. At the same time, the latter aims to design waste and pollution out of economic activity, keeping materials in use for as long as possible [51,52]. As these goals gain momentum on the global stage, there is growing recognition that traditional approaches alone will not suffice. What is required is not just incremental improvement, but transformative change that is increasingly enabled by the rapid evolution of technology, particularly artificial intelligence (AI). AI, in its various forms —machine learning, deep learning, neural networks, and hybrid systems —offers a powerful toolkit for tackling the complexities inherent in Net Zero and circular economy transitions. Unlike conventional systems that depend on fixed rules or static models, AI can learn from vast amounts of data, adapt to changing conditions, and uncover patterns and insights that might otherwise go unnoticed. This ability to process, analyse, and act upon information at scale and speed makes AI uniquely suited to address the dynamic, interconnected challenges of decarbonising economies and closing material loops [53, 54].

To understand the role of AI in achieving Net Zero, it is essential to consider the nature of the energy systems that underpin modern life. From electricity grids and transportation networks to heating systems and industrial processes, these systems are characterised by complexity, variability, and interdependence. The integration of renewable energy sources, such as wind and solar, is essential for reducing carbon emissions. However, it introduces new challenges due to its intermittent nature. Unlike fossil fuel-based generation, renewables cannot be dispatched on demand in the same way; they depend on environmental conditions that fluctuate over time and space. Balancing supply and demand in such systems requires advanced forecasting, flexible infrastructure, and intelligent control. This is where AI truly shines. By analysing historical data alongside real-time inputs from sensors, weather models, and smart meters, AI algorithms can predict energy generation and consumption patterns with remarkable accuracy. This predictive capability allows grid operators to make more informed decisions about how to dispatch energy resources, when to store surplus power, and how to manage demand peaks. For instance, machine learning models can forecast solar generation at a given site based on cloud cover data, historical irradiance, and other variables, enabling better integration of distributed solar systems into the grid. Similarly, AI-driven demand response systems can identify opportunities to shift consumption to periods of low-carbon supply, reducing reliance on carbon-intensive peaking plants. However, the role of AI extends beyond the grid. In buildings, which account for a significant share of global energy use and emissions, AI-enabled systems can optimise heating, cooling, lighting, and appliance operation in real time. By learning from occupancy patterns, weather conditions, and user preferences, these systems can maintain comfort while minimising energy use. On a larger scale, AI can support urban planners in designing energy-efficient communities by modelling how different layouts, materials, and technologies impact overall carbon performance [55, 56].

Transport, another major contributor to emissions, is also being transformed by AI. Autonomous vehicles, smart traffic management, and electrification are all areas where AI plays a central role. For example, routing algorithms can minimise congestion and associated emissions. At the same time, AI systems can optimise charging schedules for fleets of electric buses or delivery vans to align with periods of low-carbon electricity supply. The potential gains from such applications are substantial, not only in terms of carbon reduction but also in terms of improved air quality, enhanced health, and increased economic productivity. When we shift focus to the circular economy, AI's potential is equally profound. The circular economy challenges us to rethink how we design, use, and recover materials and products. Rather than following the traditional linear model of take-make-dispose, circular systems aim to keep resources in use for as long as possible, extracting maximum value and regenerating natural systems. Achieving this vision requires a deep understanding of material flows, product lifecycles, and consumer behaviour, areas where AI can provide critical insights [57, 58].

One of the key enablers of a circular economy is the ability to track and trace materials and products through complex supply chains. AI technologies, combined with the Internet of Things (IoT), can create digital twins of products, capturing data on their composition, condition, location, and history. This data can then inform decisions about maintenance, refurbishment, remanufacturing, or recycling. For example, AI can help identify which components of a used product are suitable for reuse, which require repair, and which should be recycled, thereby improving efficiency and reducing waste. AI also enhances sorting and recycling processes, which are often bottlenecks in circular systems. Traditional sorting relies heavily on manual labour or mechanical processes that can struggle with complex, mixed waste streams. AI-enabled vision systems, combined with robotics, can rapidly and accurately identify different materials, separating them for higher-quality recycling. These systems can adapt to changing input streams and improve over time as they learn from data, overcoming some of the limitations of conventional approaches. Predictive analytics is another area where AI supports circularity. By forecasting demand for refurbished or remanufactured products, businesses can better align supply with market needs, reducing overproduction and waste. Similarly, AI can help design products for longevity and disassembly, analysing data on failure modes and usage patterns to inform design choices that extend product life and facilitate end-of-life recovery. Crucially, AI can also support the business models that underpin the circular economy. Product-as-a-service models, where consumers pay for access rather than ownership, depend on the ability to monitor and manage assets in use. AI systems can optimise maintenance schedules, predict failures, and ensure that assets are used efficiently and returned at the end of their contract. This not only enhances customer experience but also supports the economic viability of circular business models. While the potential of AI is considerable, it is essential to recognise that technology is only part of the solution. Achieving Net-Zero and circular economy goals requires systemic change that encompasses policy reform, cultural shifts, and new forms of collaboration across sectors and disciplines. AI must be deployed thoughtfully, with attention to ethical considerations, inclusivity, and transparency. This is particularly important, given the risk that AI systems can inadvertently embed biases or exacerbate inequalities if not carefully designed and governed [59, 60].

One area of growing importance is explainable AI (XAI). As AI systems become more complex, there is a risk that they turn into black boxes, making decisions that are difficult for humans to understand or challenge. In sectors such as energy and waste management, where decisions can have significant environmental and social consequences, transparency is vital. XAI techniques, such as SHAP values or LIME, can help stakeholders understand how AI models arrive at their recommendations, building trust and supporting accountability.

Another critical consideration is data governance. AI models depend on data, and ensuring that this data is accurate, representative, and used responsibly is fundamental to their success. This includes addressing privacy concerns, securing data against misuse, and ensuring that data-driven insights are shared in ways that support collective action rather than entrenching proprietary advantage. The role of AI in achieving Net Zero and circular economy goals is therefore not just technical but also socio-political. It involves shaping how technologies are developed, who benefits from them, and how they are integrated into broader systems of governance and practice. For AI to fulfil its potential as an enabler of sustainability, it must be aligned with the principles of fairness, inclusivity, and environmental justice [61,62].

The PhD research that underpins this discussion takes these complexities seriously. It does not treat AI as a magic bullet but as a tool that, when carefully designed and thoughtfully applied, can support the transition to more sustainable energy and material systems. Through the development of predictive models for energy demand, recycling rates, and waste generation, the work demonstrates how AI can improve accuracy, efficiency, and responsiveness in key areas. By integrating explainable AI techniques, it addresses the need for transparency and stakeholder confidence. By designing flexible and adaptable models, it acknowledges the diversity of contexts in which Net Zero and circular economy strategies must be implemented. In sum, AI holds significant promise as a driver of the transformative change needed to achieve Net Zero and circular economy goals. Its strength lies in its ability to handle complexity, adapt to changing conditions, and generate insights that support smarter, more sustainable decisions. But realising this promise will require more than technical ingenuity. It will demand a commitment to ethical, inclusive design; robust governance; and a willingness to place technology in the service of broader societal and ecological objectives. The path to Net Zero and a circular economy is not easy. However, with the thoughtful integration of AI, it becomes a more navigable and attainable path within our collective reach [63].

1.4 Motivation and Rationale for the Study

The journey towards sustainability is one of the defining challenges of our time. Across the globe, nations, industries, and communities face the pressing need to reconcile economic development with the urgent demands of environmental stewardship. Nowhere is this challenge more acute than in the sectors of energy and waste management. These systems, which are vital to modern life, are also among the most significant contributors to environmental degradation and climate change. Against this backdrop, the motivation for this study stems from recognising the scale and complexity of these problems and the conviction that innovative solutions, particularly those leveraging advanced computational tools, are necessary to chart a more sustainable path forward.

To begin with, the motivation for this research is deeply rooted in the realities of the climate crisis. The science is unequivocal: human activities, especially the burning of fossil fuels and the mismanagement of resources, are driving global temperatures upwards, with far-reaching and often devastating consequences. From rising sea levels that threaten coastal communities, to shifting weather patterns that endanger food security, to the loss of biodiversity that undermines the resilience of ecosystems, the impacts are visible and growing. In this context, achieving Net Zero greenhouse gas emissions by mid-century has become not just an aspirational target, but an imperative.

However, translating the Net Zero goal into reality is no simple task. Energy systems, which lie at the heart of the carbon challenge, are complex, interconnected, and often inflexible. Many existing infrastructures were designed in an era when fossil fuels were abundant, and the environmental costs of carbon emissions were not fully understood or considered. Today, as the world moves toward integrating renewable energy sources, such as solar and wind, that are inherently variable, the limitations of traditional energy management approaches become increasingly apparent. Balancing supply and demand, ensuring grid stability, and optimising consumption patterns in this new landscape demand a level of intelligence and adaptability that conventional rule-based systems struggle to provide.

Parallel to the energy challenge is the crisis of waste. The world generates billions of tonnes of waste every year, much of which ends up in landfills or, worse, in the natural environment.

Plastics pollute the oceans, electronic waste leaches toxins into the soil, and food waste contributes to unnecessary greenhouse gas emissions. The circular economy offers a compelling alternative in which materials are kept in use for as long as possible, waste is designed out of systems, and natural systems are regenerated. Nevertheless, making the circular economy a practical reality, particularly at scale, requires an ability to track materials, predict flows, and optimise processes in ways that are simply beyond the capabilities of manual planning or static models. The twin challenges of energy and waste, along with the transformative potential of the Net Zero and circular economy frameworks, form the backdrop against which this study is set. The rationale for focusing on artificial intelligence (AI) is based on the belief, supported by emerging evidence and technological progress, that AI offers unique capabilities for addressing the complexities of these systems. AI can process vast amounts of data from diverse sources, learn patterns that are too subtle or complex for human analysts to detect, and support decision-making that is both more precise and more adaptive.

Importantly, this study is motivated not just by the promise of AI in the abstract, but by a recognition of specific gaps and limitations in current practice. In energy management, many existing systems rely on fixed schedules or simple heuristics that do not account for the dynamic interplay of factors such as weather conditions, occupancy patterns, and energy prices. Demand response programmes, for example, often use broad signals or time-of-use pricing to shift consumption, but without the necessary granularity or predictive accuracy to fully realise efficiency gains or reduce emissions. Similarly, in waste management, collection schedules are frequently static, failing to respond to actual variations in waste generation across neighbourhoods or over time. Recycling targets are set, but the tools for achieving them often lack the data-driven insights needed to optimise processes on the ground. The study is driven by the conviction that computational artificial intelligence, encompassing machine learning, deep learning, and hybrid models, can bridge these gaps. By harnessing data from smart meters, IoT sensors, weather stations, and waste collection records, AI models can provide forecasts and recommendations that are tailored, responsive, and far more accurate than traditional approaches. For instance, machine learning models can predict residential energy demand at fine temporal and spatial scales, enabling more targeted demand response and better integration of renewable sources. Long short-term memory (LSTM) networks, a type of recurrent neural network well-suited for time-series data, can capture temporal dependencies in recycling rates, thereby supporting more effective planning of waste collection and processing.

A further source of motivation is the need for transparency and trust in these systems. As AI models become more complex, there is a risk that they become black boxes, making predictions or recommendations that are difficult for human operators to understand or validate. This study recognises that for AI to be adopted in critical infrastructure systems, such as energy and waste management, it must be explainable. People need to understand not just what the models predict, but why they make those predictions. To this end, the research integrates explainable AI techniques, such as SHAP and LIME, which provide insights into how input variables influence model outputs. This emphasis on explainability is driven by a commitment to responsible and ethical AI, as well as the practical need to build confidence among decision-makers, regulators, and the public. Another important strand of motivation stems from recognising the diversity of contexts in which energy and waste challenges are played out. The solutions developed in this research are designed to be flexible and adaptable, capable of being applied in different geographic regions, climatic conditions, and socio-economic settings. While the empirical work focuses on particular case studies, such as residential energy demand or municipal recycling rates, the underlying modelling frameworks are intended to be generalisable. This reflects an awareness that sustainability challenges and opportunities vary across contexts, and that scalable solutions must be sensitive to local conditions.

The rationale for this study is also shaped by the urgent timelines involved. The window for meaningful action to address climate change is narrow. According to the Intergovernmental Panel on Climate Change (IPCC), global emissions must be halved by 2030 to limit warming to 1.5°C and avoid the most severe impacts of climate change. This urgency compels the search for solutions that can be deployed rapidly and effectively. AI, with its capacity for rapid learning and deployment, offers a promising avenue in this regard. By demonstrating how predictive models can enhance the efficiency and sustainability of energy and waste systems, this research aims to provide practical tools that support immediate and impactful action. At the same time, the study is motivated by an awareness of the risks and challenges associated with AI. These include issues of data privacy, cybersecurity, potential biases in training data, and the socio-economic impacts of automation. The research takes these challenges seriously, embedding considerations of ethics and governance into the development and evaluation of models. This reflects a broader view that technological innovation must be aligned with societal values and designed to serve the public good.

In addition to the environmental and technical drivers, the study is motivated by policy and market trends. Governments around the world are increasingly committing to Net Zero targets, while industries are seeking to align with circular economy principles as a source of competitive advantage and resilience. Yet, there is a gap between ambition and implementation. This research aims to help bridge that gap by providing models and frameworks that can support evidence-based decision-making, operational optimisation, and continuous learning. The rationale for combining energy and waste in a single study stems from their interconnections. Energy is used throughout the waste management chain, from collection and transport to processing and disposal. Waste, in turn, can be a source of energy through technologies such as anaerobic digestion, gasification, and waste-to-energy incineration. By developing integrated models that consider these linkages, the study aims to support more holistic and synergistic approaches to sustainability.

Finally, the personal motivation for the study should not be overlooked. The researcher's commitment to contributing to a more sustainable and equitable world underpins the work. The study reflects a belief in the power of science and technology to drive positive change, while also acknowledging that such change requires thoughtful design, collaboration across disciplines, and engagement with diverse stakeholders. It is driven by curiosity, a desire to solve complex problems, and a sense of responsibility towards future generations. In conclusion, the motivation and rationale for this study are multifaceted. They arise from the scale and urgency of global sustainability challenges, the potential of AI to provide transformative solutions, the gaps and limitations of current practices, and a commitment to ethical, inclusive, and impactful research. By advancing our understanding of how computational intelligence can support Net Zero and circular economy goals, the study aims to make a meaningful contribution to one of the most critical missions of our time. Figure 1 highlights the key limitations in sustainable energy and waste systems management that necessitated the creation of AI-driven predictive frameworks in this study. It highlights the disintegration of data, inadequate forecasting instruments, and the absence of cohesive decision-support systems, prompting the adoption of AI to address these issues and facilitate more intelligent sustainability initiatives.

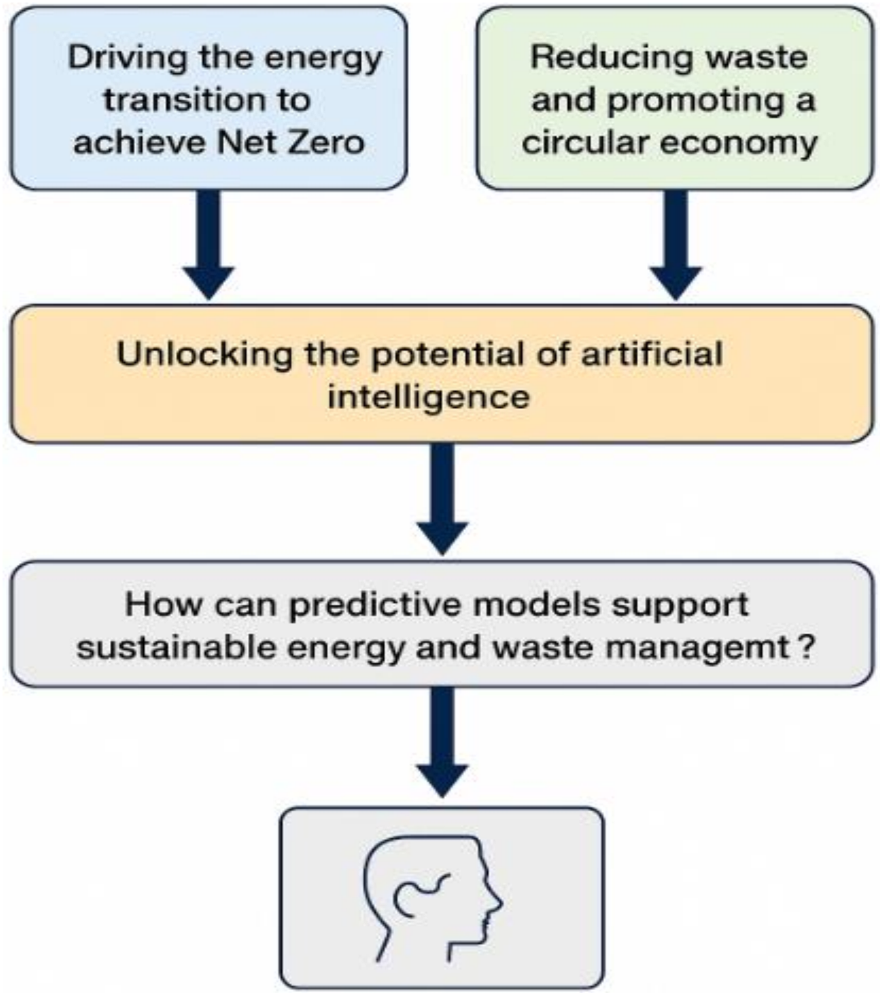

Figure 1.2. Rationale for this research and unresolved questions before publication.

1.5 Research Aims and Scope

This research aims to develop a novel, computationally intelligent framework that integrates predictive modelling of energy consumption and waste generation to advance sustainable practices within the circular economy. Specifically, the study investigates the application of artificial intelligence (AI), including machine learning (ML), deep learning (DL), and explainable AI (XAI), to optimise decision-making processes in energy and waste systems at municipal and industrial levels. The scope of this PhD spans multiple interconnected domains: renewable energy forecasting, waste classification and reduction modelling, and the development of integrated platforms capable of delivering actionable insights for policy and infrastructure design. While traditional studies focus on either energy or waste in isolation, this research fuses both systems into a holistic predictive architecture, enabling real-time data analysis, resource optimisation, and sustainability tracking. The study is bounded by practical considerations such as data availability, computational capacity, and region-specific environmental variables. Case studies and simulation environments are used to validate models. In sum, this research sets out not only to innovate technically but also to offer a scalable, interdisciplinary approach to achieving net-zero and circular economy goals through the power of AI.

1.6 Specific Research Objectives

1. To develop and validate AI-based models for forecasting short- and long-term energy demand patterns.
2. To construct intelligent models for predicting and classifying waste generation by type, volume, and source.
3. To design a unified AI framework that simultaneously manages energy and waste data for circular economy decision-making.
4. To incorporate explainable AI (XAI) and user-focused dashboards into the system for transparency, accountability, and practical adoption.
5. To evaluate the performance, scalability, and transferability of the integrated system through real-world case studies or simulations.

1.7 Research Questions Hypothesis

1 How can artificial intelligence techniques be effectively employed to predict and optimise energy consumption across various sectors?

2 What models can be developed to accurately predict waste generation patterns based on socio-economic, demographic, and environmental indicators?

3 Can an integrated predictive framework be designed to manage both energy and waste systems simultaneously under circular economy principles?

4 How can explainable AI (XAI) tools be applied within this system to ensure transparency, trust, and real-world usability by non-expert stakeholders?

1.7 Significance of the Study

This study offers a timely and innovative contribution to addressing two of the most pressing challenges of our time: sustainable energy management and waste reduction by harnessing the transformative power of artificial intelligence (AI). Through the integration of machine learning models such as LSTM, MLP, Gradient Boosting, and Random Forest, the research presents a data-driven framework capable of forecasting, optimising, and supporting decision-making processes in complex environmental systems. The ability to extract predictive insights from time-series data (e.g., smart meters, waste logs, and climate sensors) opens new avenues for improving policy, infrastructure design, and resource allocation across sectors.

Crucially, the study aligns with the United Nations Sustainable Development Goals (SDGs), particularly those targeting clean energy (SDG 7), responsible consumption (SDG 12), and climate action (SDG 13). Furthermore, the integration of Explainable AI (e.g., SHAP, LIME) enhances trust, transparency, and accountability in model deployment, paving the way for real-world adoption by municipalities, utility providers, and environmental agencies. By bridging the gap between theoretical modelling and practical application, this research not only advances academic discourse but also offers measurable societal and environmental impact, reaffirming the role of AI as a critical enabler of circular economy transitions and net-zero ambitions.

1.8 Scope of the Study

This study focuses on the development and application of AI-driven predictive models, specifically LSTM, MLP, Gradient Boosting, and Random Forest for optimising energy consumption and municipal waste management systems. The scope includes data sourced from smart meters, sensors, weather logs, and waste collection records across selected urban regions. It also encompasses the integration of Explainable AI (SHAP, LIME) to enhance the transparency and interpretability of the models. By addressing forecasting, resource optimisation, and model explainability, the study targets scalable, sustainable solutions within the broader context of circular economy goals and net-zero transitions in environmental infrastructure.

1.9 Structure of the thesis

Chapter 1: Introduction - The chapter begins with a discussion of global sustainability challenges, focusing in particular on the energy and waste management sectors. Research Aims, Objectives, and Questions - This builds directly on the groundwork laid in the introduction and literature review. It presents a detailed articulation of the research aims and objectives, ensuring that these are linked to the identified gaps and needs outlined earlier. The chapter also sets out the specific research questions that guide the study, explaining how these questions are designed to elicit new insights and contribute original knowledge.

Chapter 2: Literature Review - The chapter provides a comprehensive review of existing scholarship and practice in the relevant fields.

Chapter 3: Methodological Framework - The fourth chapter moves the thesis to its technical core. It presents the methodological choices made in the study, explaining the rationale for selecting particular computational techniques, data sources, and analytical procedures. The chapter begins with a discussion of the research philosophy that underpins the work, reflecting on questions of epistemology, validity, and reliability.

The subsequent sections describe the design and development of the predictive models used in this research. This includes an explanation of how data from various sources, such as smart meters, weather sensors, and municipal records, was collected, cleaned, and processed. The chapter details the structure and parameters of the machine learning and deep learning models employed, including multilayer perceptrons (MLPs), long short-term memory (LSTM) networks, and ensemble approaches. Attention is also given to the use of explainable AI techniques, such as SHAP and LIME, and the reasons for integrating these into the model development process.

Chapter 4 focuses on the first publication, which addresses the development of AI-based predictive models for residential energy demand response. This chapter situates the publication into the broader research aims, summarises its key findings, and offers reflections on its methodological and practical significance.

Chapter 5 presents the second publication, which explores the use of LSTM networks in predicting municipal recycling rates. The chapter highlights the novel aspects of this work, such as its treatment of temporal dependencies and its contribution to circular economy planning.

Chapter 6 covers the third publication, which deals with a comparative performance analysis of rule-based versus AI-driven models in energy demand management. It draws out the implications of these findings for policy and practice, particularly in the context of achieving Net Zero goals.

Chapter 7: Synthesis of Findings and Contribution to Knowledge - Chapter 8 serves the essential function of drawing these together into a cohesive whole. It synthesises the key findings from the various studies, highlighting the common threads, synergies, and cumulative insights that emerge.

Chapter 8: Conclusion and Future Directions - The final chapter offers a closing reflection on the research journey. It begins by summarising the main findings and contributions, providing a clear and concise restatement of how the thesis has advanced knowledge and practice.

Chapter 2

2 Literature Review

2.1 Overview of Energy and Waste Management Systems.

Energy and waste management systems are central to sustainable development strategies, particularly in the context of climate change, urbanisation, and global resource scarcity. Modern energy systems are evolving from centralised, fossil fuel-based networks to decentralised smart grids incorporating renewable sources such as solar and wind. These new systems demand real-time data analysis, predictive capabilities, and load-balancing mechanisms that traditional infrastructure cannot adequately support [64, 65]. Similarly, waste management is shifting from a linear model of collection and disposal to circular models that prioritise resource recovery and closed-loop systems [66, 67]. Recent studies have highlighted the limitations of conventional energy forecasting models in the face of high variability and intermittency of renewables. Traditional statistical approaches, while useful for linear patterns, often fail to capture nonlinear dependencies and complex temporal trends associated with smart grid operations [68, 69]. Researchers have increasingly turned to artificial intelligence (AI), particularly machine learning (ML) and deep learning (DL) techniques, to improve forecasting accuracy for both energy demand and renewable supply [70,71]. These models demonstrate superior performance in recognising consumption patterns, integrating weather variables, and adapting to dynamic behaviours of users and systems [72].

In parallel, the field of waste management has begun to embrace computational methods, though progress remains slower. Most municipal waste management systems still operate on static routing and collection schedules, which contribute to inefficiencies, fuel wastage, and environmental impacts [73]. While geographic information systems (GIS) and optimisation algorithms have been explored to some extent, integration with AI remains limited [74]. Moreover, few studies consider the interdependence between energy and waste systems, for instance, the potential of waste-to-energy conversion or the energy footprint of recycling processes [75]. A significant body of literature acknowledges the potential of AI to support sustainable waste strategies through predictive modelling, intelligent bin-level sensors, and robotic waste sorting [76], [77]. However, these studies tend to focus narrowly on isolated tasks rather than building end-to-end decision support systems that are scalable and explainable [78]. Moreover, the application of hybrid AI models that combine traditional forecasting with real-time data integration is still underexplored in both sectors [79].

One of the most critical gaps in the existing literature is the absence of integrated frameworks that simultaneously optimise energy and waste systems using AI. Despite shared data streams (e.g., from IoT devices), overlapping stakeholders, and common sustainability targets, very few studies bridge these domains in a unified computational model [80], [81]. Additionally, the lack of explainable AI (XAI) in most applied research raises concerns about transparency, stakeholder trust, and ethical deployment [82], [83]. This PhD research addresses these gaps by proposing AI-driven predictive models that not only improve operational efficiency in energy and waste systems but also provide explainable outputs to enhance trust and adoption. Figure 3 shows the holistic view that emphasises the integration potential of waste and energy systems in achieving net-zero emissions and circular economy goals.

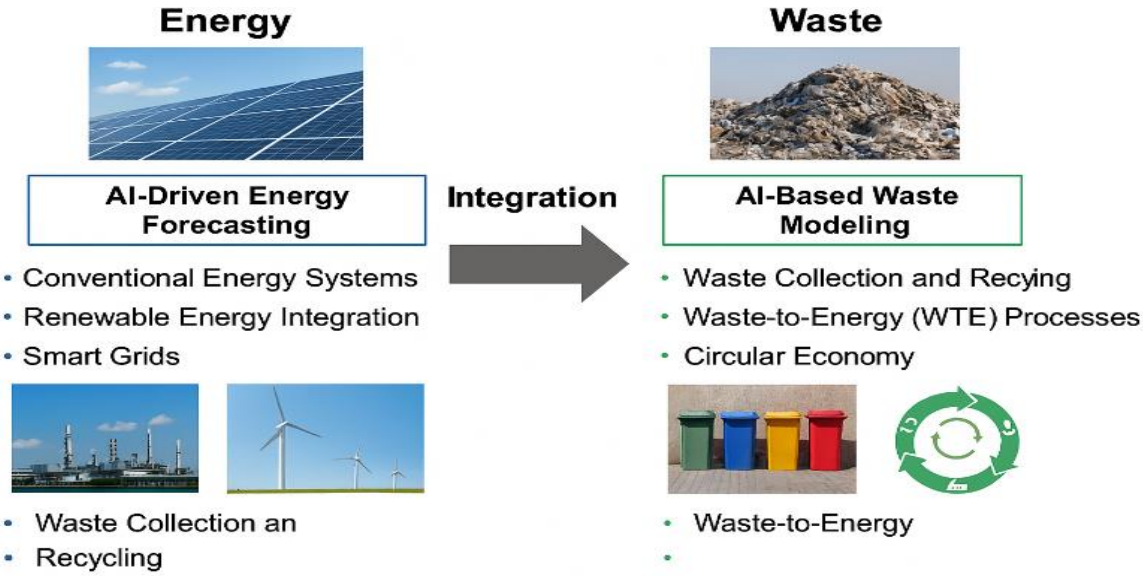

Figure 2.1. The overview of energy and waste management systems

2.2 Predictive Analytics and Computational Intelligence in Sustainability

The growing urgency of climate change and the quest for sustainable development have significantly amplified interest in data-driven solutions that support predictive decision-making in environmental systems. Predictive analytics and computational intelligence, especially methods like machine learning (ML), deep learning (DL), and hybrid neural networks, have emerged as powerful tools in modelling complex systems such as energy consumption, carbon emissions, and waste flows [84, 85]. These techniques offer capabilities well beyond traditional statistical models by enabling the discovery of nonlinear patterns and real-time system adaptations in highly variable data environments [86].

Recent studies have applied artificial intelligence (AI) to forecast electricity demand at micro and macro levels, particularly in smart grid contexts. For instance, models such as multilayer perceptrons (MLPs), support vector machines (SVMs), and long short-term memory (LSTM) networks have shown strong performance in managing the stochastic nature of solar radiation, energy pricing, and user consumption patterns [87, 88, 89]. These models have enabled grid operators to optimise resource allocation and reduce peak loads, contributing directly to energy efficiency and climate mitigation efforts.

Parallel advancements in waste management modelling have been made, particularly in optimising collection routes and forecasting municipal solid waste generation [90, 91]. However, the depth and sophistication of computational intelligence in waste systems remain significantly behind those applied in energy sectors. While smart bins and IoT sensors have enabled some reactive data collection, there is limited predictive integration that utilises real-time streams to anticipate waste surges or facilitate automated material recovery strategies [92, 93]. Furthermore, the literature shows that most current applications of predictive analytics in sustainability remain narrowly focused on single-issue optimisation, such as forecasting electricity demand or route planning for waste vehicles. Few efforts have explored the broader potential of these technologies to inform strategic sustainability planning, such as policy impact forecasting or lifecycle optimisation [94, 95]. Additionally, while computational intelligence methods like genetic algorithms and fuzzy logic have been applied in specific sustainability use-cases, they are often deployed in isolation rather than integrated into cohesive decision-support frameworks that combine forecasting with prescriptive analytics [96, 97].

Another underdeveloped area is explainable artificial intelligence (XAI) within sustainability. Many high-performing models are ‘black boxes,’ offering limited insight into how predictions are formed. This issue undermines trust and limits practical adoption by urban planners, engineers, and sustainability officers [98], [99]. Although frameworks such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) have been introduced in the AI field, they are yet to be widely adopted in sustainability domains, especially in low-resource settings [100], [101].

The evident gap in the literature lies in the absence of unified, explainable, and scalable predictive systems that connect energy and waste domains using real-time computational intelligence. There is also a lack of emphasis on hybrid models that integrate physical simulation, machine learning, and policy scenarios to improve both operational and strategic decision-making in sustainability systems [102, 103]. This PhD study addresses these gaps by developing AI-driven predictive models that are not only accurate but also interpretable and capable of cross-sectoral integration, advancing both the theoretical and practical frontiers of sustainable systems management.

2.3 Traditional Rule-Based Methods vs. Machine Learning Approaches.

Rule-based systems have long been employed in decision-making processes related to energy and waste management. These systems typically rely on expert-defined conditions, thresholds, and deterministic logic to control system behaviours or guide decisions. In their simplest form, rule-based models provide transparency and ease of implementation, making them ideal for scenarios where system dynamics are well-understood and remain relatively stable over time [104,105]. For instance, in early smart grid operations, threshold-based rules were used to trigger load shedding, initiate demand response events, or reroute energy flow in the case of faults [106].

However, as energy and waste systems have grown more complex —driven by rising demand, the integration of renewable sources, and urban population growth —rule-based systems have shown serious limitations. These approaches struggle to adapt to environments characterised by uncertainty, nonlinearity, and real-time feedback loops [107]. In the context of municipal waste systems, for example, rule-based routing algorithms often fail to optimise collection schedules effectively when bin fill levels vary daily or when unexpected road closures and traffic conditions interfere with predefined plans [108, 109]. Similarly, in energy forecasting tasks, fixed rules cannot capture fluctuations in user behaviour, weather patterns, or distributed energy resource outputs [110].

To address these shortcomings, researchers have increasingly turned to machine learning (ML) and artificial intelligence (AI) as alternative approaches. ML techniques, unlike rule-based systems, learn from historical and real-time data to discover hidden patterns, classify system states, and predict future outcomes [111, 112]. In smart energy systems, ML has been used to forecast load demand, predict photovoltaic (PV) output, and optimise dynamic pricing models more accurately than traditional methods [113, 114]. Deep learning models, particularly Long Short-Term Memory (LSTM) networks and convolutional neural networks (CNNs), have shown remarkable accuracy in capturing temporal dependencies in energy time-series data [115, 116].

The transition to ML methods has also begun in waste management domains. Intelligent models now predict daily or weekly waste generation volumes, classify recyclable materials, and even power robotic waste sorting systems [117, 118]. These applications highlight the adaptability of ML in contexts with high variability and incomplete data conditions under which rule-based methods typically fail [119]. Notably, recent studies have also shown the use of hybrid approaches, where rule-based constraints are embedded within ML frameworks to preserve safety boundaries or regulatory requirements [120].

Despite the advantages of ML, challenges remain, especially concerning model interpretability and deployment in real-world systems. While rule-based systems offer simplicity and transparency, ML models are often perceived as “black boxes,” making it difficult for operators, policymakers, and stakeholders to understand or trust their outputs [121, 122]. This issue is particularly pressing in safety-critical domains like waste treatment facilities or energy grid control centres, where the consequences of wrong predictions could be severe [123].

Furthermore, rule-based models still dominate in regions or sectors with limited access to data or computational resources. In many low-income municipalities, for example, real-time data infrastructure is either missing or unreliable, making it difficult to train or validate complex ML models [124]. This has led to a dual-system landscape, where rule-based and ML systems coexist, often with limited integration. There is a growing consensus that the future lies in developing explainable, data-efficient ML systems that can complement or gradually replace rigid rule-based logic where appropriate [125], [126].

The key gap in the literature lies in the lack of unified frameworks that allow for seamless transition or integration between rule-based and ML-driven decision systems. Most current research either focuses on ML's superiority in accuracy or critiques its lack of interpretability. However, very few studies propose structured methodologies for combining the robustness of rules with the flexibility of learning systems [127]. Moreover, comparative performance evaluations under identical operational scenarios, especially across multiple domains such as energy forecasting, waste routing, and material recovery, are scarce. This hampers a fair assessment of trade-offs between the two paradigms in real-world deployments [128]. This PhD research addresses these gaps by developing and evaluating AI-driven predictive models for energy and waste systems, benchmarking them against traditional rule-based approaches. It also integrates explainability modules and hybrid architecture considerations, thereby contributing both technically and methodologically to the current body of knowledge. Figure 4 shows why rule-based systems are insufficient for complex, sustainability-focused infrastructure and how ML approaches enhance decision-making, prediction accuracy, and system adaptability in energy and waste systems.

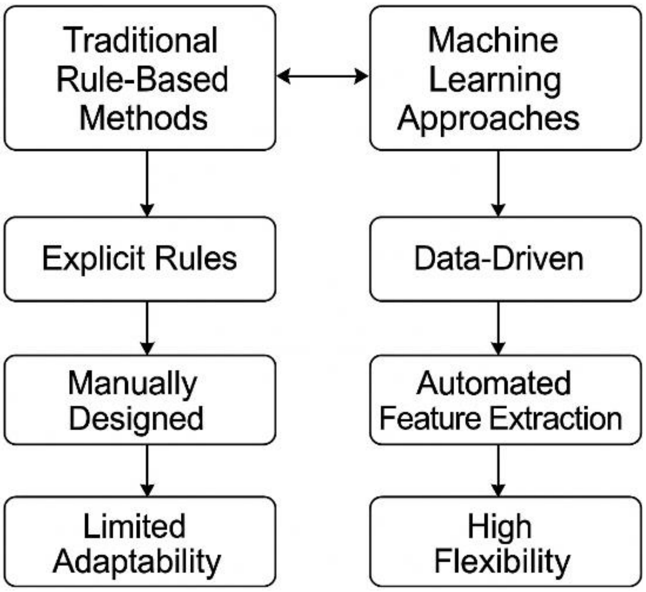

Figure 2.2 Comparison between traditional rule-based methods and machine learning approaches.

2.4 Deep Learning (e.g., LSTM) and Its Application in Waste and Energy Systems.

Deep learning (DL) has emerged as one of the most transformative innovations in the modelling of complex systems over the past decade, especially in the fields of energy and waste management. Among various DL architectures, Long Short-Term Memory (LSTM) networks have gained prominence due to their ability to learn long-range temporal dependencies and nonlinear relationships in time-series data. Initially designed for sequence modelling tasks such as language processing, LSTM networks have since been adapted for energy load forecasting, renewable energy generation prediction, and even waste volume estimation [129, 130].

In energy systems, traditional forecasting techniques often rely on autoregressive integrated moving average (ARIMA), regression models, or rule-based estimations that fail to capture the volatility and dynamic behaviour of smart grid environments [131, 132]. By contrast, LSTM models learn from historical trends and can adapt to seasonal patterns, user consumption fluctuations, and weather-driven variability in renewable sources like wind and solar [133, 134]. For instance, Marino et al. demonstrated the superiority of LSTM in predicting hourly energy demand in urban grids, outperforming feed-forward neural networks and statistical baselines in both accuracy and robustness [135]. Other researchers have extended the application of LSTM to real-time load dispatch, anomaly detection, and energy price prediction. Kong et al. developed a hybrid LSTM model integrated with a convolutional neural network (CNN) to improve short-term load forecasting in smart grids, showing that combining spatial and temporal features significantly enhances prediction precision [136]. Similarly, Wang et al. used an LSTM-Attention model to forecast solar energy output, achieving notable improvements in regions with high meteorological variance [137, 138]. These models enable utilities to plan better energy storage usage, shift loads, and mitigate system instabilities. While energy systems have seen an abundance of LSTM-based research, waste management applications have only recently begun to explore deep learning methods. Historically, waste systems were modelled using static data inputs or simplistic projections based on population and urban density [139, 140]. However, the growth of smart cities and IoT-enabled infrastructure has made it possible to collect high-resolution waste generation data, enabling the training of LSTM models that can predict bin fill levels, waste generation volumes, and even recycling behaviour patterns [141, 142].

Several pilot studies have explored LSTM for predicting municipal solid waste generation at city or district scales. In one such study, Chowdhury et al. applied LSTM to forecast daily waste generation in Dhaka, Bangladesh, using a combination of weather, event, and socio-economic data [143]. The model showed strong predictive power and offered a path for integrating waste forecasting into smart bin logistics. However, many of these studies remain limited to proof-of-concept and are not yet integrated into city-wide waste systems. Furthermore, a significant gap exists in the use of LSTM for process-level waste systems, such as recycling stream optimisation, landfill gas emissions prediction, or waste-to-energy plant operations. The complexity of these systems due to multiple input streams, delayed feedback loops, and regulatory constraints makes them ideal candidates for recurrent neural networks like LSTM. Yet, few studies have explored this potential [144, 145]. For example, predicting the methane output of landfills using historical temperature and organic matter data could help in energy recovery planning, but this remains largely unexplored in academic literature. Another challenge is model interpretability. LSTM networks, despite their predictive accuracy, are often considered “black boxes,” limiting their adoption in high-stakes decision environments like energy control centres or waste treatment facilities. Few studies incorporate explainable AI (XAI) techniques to make the internal logic of LSTM models transparent and auditable for end-users and stakeholders [146, 147]. This is especially critical in policy-driven sectors like environmental management, where accountability and regulatory compliance are essential.

Moreover, hybrid modelling techniques where LSTM is combined with physics-based simulations, reinforcement learning, or optimisation heuristics remain underutilised in both energy and waste domains. While some progress has been made in hybridising DL with control algorithms for microgrid management, such practices are rarely extended to closed-loop waste systems or circular economy scenarios [148, [149]. In summary, while LSTM has proven to be a robust tool in energy demand forecasting and solar power prediction, its application in waste management is still at a nascent stage. The key gaps in the literature include (1) limited deployment of LSTM in waste systems beyond basic forecasting, (2) insufficient integration of explainability features, and (3) a lack of hybrid, cross-domain models that link waste and energy systems holistically. This PhD research aims to bridge these gaps by developing integrated, interpretable LSTM-based frameworks that can support predictive planning across both sectors, offering a scalable and actionable contribution to sustainable system design. Figure 5 provides a visual representation of how Long Short-Term Memory (LSTM) and other deep learning models operate and are applied across the energy and waste management sectors.

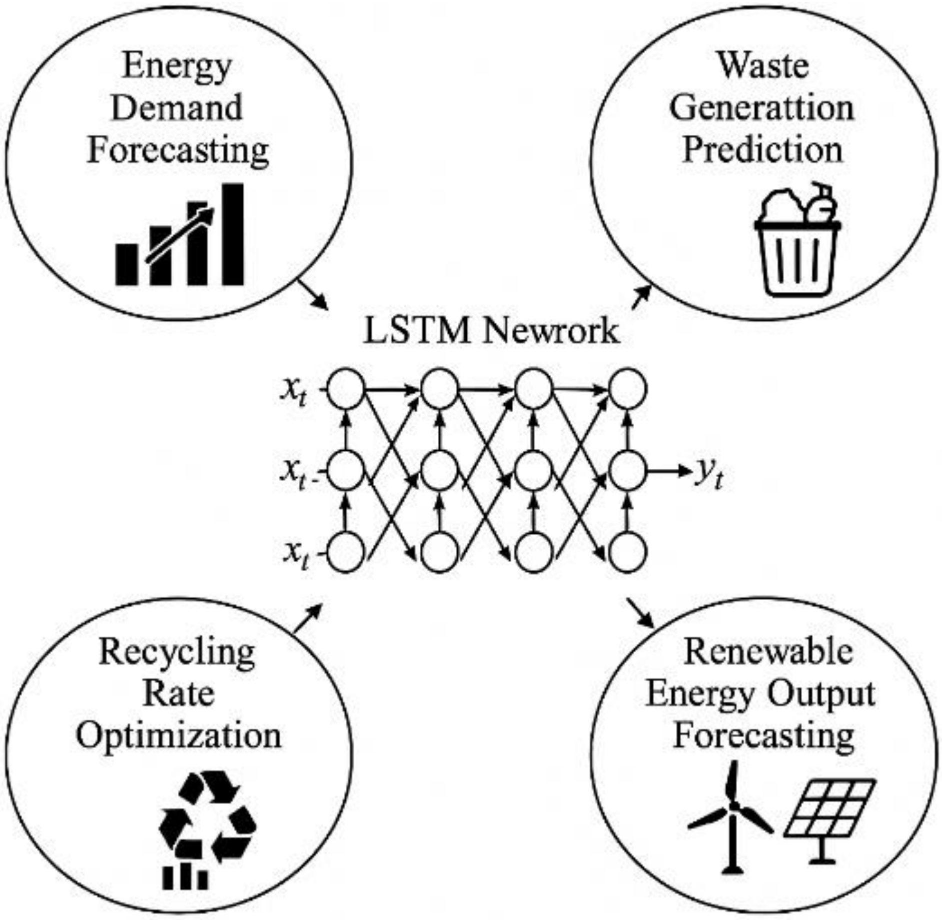

Figure 2.3 The deep learning (e.g., lstm) and its application in waste and energy systems.

2.5 Explainable AI (XAI): Principles and Relevance to Sustainability

The rapid integration of artificial intelligence (AI) into sustainability-related domains has brought immense capabilities in forecasting, optimisation, and decision-making. However, with the rising complexity and autonomy of these AI models, the demand for interpretability and transparency has grown significantly, particularly in high-impact fields such as energy distribution, waste management, and environmental policy. This has led to the emergence of Explainable AI (XAI), a field dedicated to making the inner workings of machine learning algorithms understandable to human users [150, 151]. Explainable AI aims to bridge the gap between the decision logic of black-box models, especially deep neural networks and the accountability standards required in regulated sectors. In sustainability contexts, decisions often involve significant economic, environmental, and social trade-offs, necessitating trust in model outputs. XAI tools such as SHAP (Shapley Additive Explanations), LIME (Local Interpretable Model-agnostic Explanations), and counterfactual explanations have thus become increasingly vital [152, 153, 154]. These tools enable stakeholders, from city planners to environmental regulators, to assess how and why an AI system arrives at a particular recommendation or prediction.

In energy systems, for example, deep learning models are frequently used to forecast demand, optimise load balancing, or manage renewable integration. However, the lack of interpretability in these models can lead to resistance in adoption, especially when operational reliability is at stake [155, 156]. Recent studies have shown that integrating XAI into energy forecasting frameworks improves user confidence and facilitates better human-AI collaboration. Ribeiro et al. [157] illustrated how model interpretability could aid operators in refining predictions for photovoltaic output under dynamic weather conditions, reducing operational risks.

In waste management, although data-driven optimisation is gaining ground, explainability remains a barrier to wider implementation. AI tools have been employed for route optimisation, recycling classification, and predictive waste generation. Yet, many municipal authorities are hesitant to adopt them due to opacity in model logic [158, 159]. For instance, when a smart system predicts a 30% surge in organic waste and reallocates bin collection frequency, stakeholders must understand which variables, such as weather, local events, and demographic shifts, led to that forecast. Without this clarity, decisions might be contested, especially in publicly accountable institutions.

Despite the evident need, XAI remains under-researched within sustainability domains. A substantial portion of existing literature focuses either on general applications of XAI in health and finance or on the technical development of explanation methods themselves [160, 161]. Very few studies examine how these techniques can be contextualised for complex, multistakeholder decision environments such as circular economy planning or carbon accounting [162]. Moreover, little effort has been made to develop domain-specific XAI frameworks that consider the unique dynamics and ethical imperatives of sustainability systems [163]. The gap in the literature is thus twofold: first, a lack of real-world case studies showing successful deployment of XAI in waste and energy systems, and second, limited research on co-designing explainable models with stakeholders to ensure relevance, usability, and trust. This PhD study seeks to address this gap by building and validating interpretable machine learning models specifically tailored for energy forecasting and waste stream analytics. These models will not only prioritise prediction accuracy but also offer transparent, actionable insights for decision-makers committed to achieving net-zero and circular economy targets. Figure 6 illustrates how XAI techniques are applied across various sustainability domains to enhance transparency, trust, and informed decision-making.

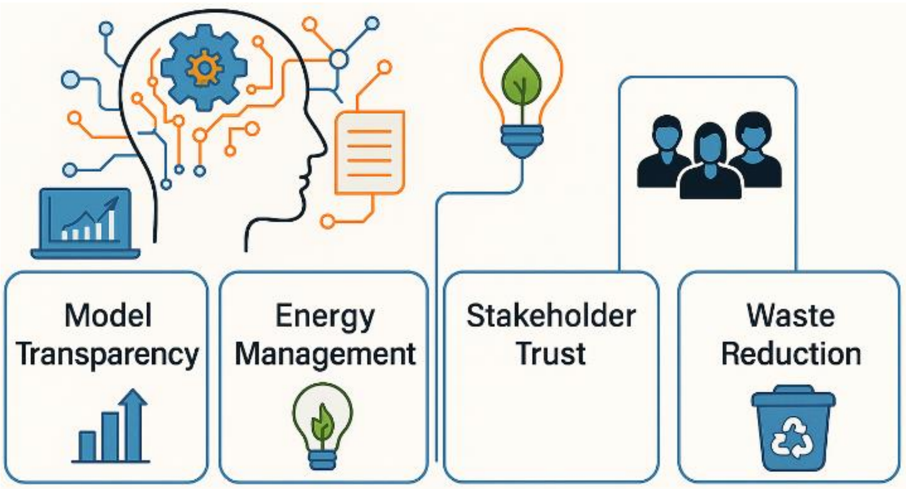

Figure 2.4 Current application of explainable ai (xai): principles and relevance to sustainability.

Chapter 3

3 Research Methodology

3.1 Research Philosophy and Design

The complexity of global sustainability challenges, particularly in the domains of energy consumption and waste management, demands a research approach that is both scientifically robust and practically responsive. As this PhD seeks to develop an integrated predictive framework driven by artificial intelligence (AI), the research philosophy and design must strike a balance between computational rigour, empirical relevance, and interpretative clarity. This chapter sets out the philosophical orientation that underpins the study, outlines the methodological framework, and explains how these choices support the research objectives and questions.

3.1.1 Research Philosophy: Pragmatic Positivism

In the field of engineering and applied artificial intelligence, research philosophies often fall within the spectrum of positivism and pragmatism. Positivism assumes that knowledge can be observed, measured, and quantified objectively, free from the influence of the researcher. This tradition is deeply rooted in the natural sciences, and it aligns with the mathematical and statistical foundations of machine learning and data modelling. However, strict positivism may overlook contextual, real-world complexity, especially in domains like waste systems and energy behaviour, which are shaped by human actions, policies, and economic incentives.

This research adopts a pragmatic-positivist stance, combining the strengths of objective analysis (through quantitative models) with an openness to the situational and interpretive realities of system design and deployment. Pragmatism, as advanced by thinkers like John

Dewey and William James, advocates for using whatever methods and theories are most effective to address a given problem. In this context, the problem is not purely theoretical; it is applied, urgent, and rooted in sustainability. Hence, while the predictive algorithms and modelling tools used in this study are grounded in positivist logic, the overarching approach is guided by pragmatism: the models must be actionable, usable by stakeholders, and contextually relevant.

A pragmatic lens allows the study to include real-world factors that may not be captured in the data alone, such as regulatory shifts, socio-economic disparities, and behavioural feedback. This hybrid philosophy also justifies the inclusion of explainable AI (XAI) tools in the study, which address not only statistical performance but also ethical, operational, and communicative considerations.

3.1.2 Research Design Overview

The research design reflects the layered complexity of the topic. The core of the study is quantitative, involving machine learning model development, data pre-processing, algorithmic optimisation, and system simulation. These tasks require a structured, iterative approach to model training and evaluation using real-world datasets. However, the design also includes qualitative reasoning where appropriate, particularly in the interpretation of results, evaluation of system scalability, and integration of stakeholder feedback into the framework's interface and usability. The study follows a deductive approach, whereby hypotheses are formulated based on the literature and tested through empirical data. For example, one assumption is that LSTM neural networks will outperform ARIMA models in forecasting non-linear energy demand patterns. Another is that computer vision algorithms can classify waste more accurately when combined with contextual metadata (e.g., source location or seasonality). These assumptions are not taken as givens; they are tested, validated, and iteratively refined throughout the research process. The research design is also applied and developmental in nature. It does not stop at generating knowledge or patterns; it proceeds to build a working predictive system that governments, municipalities, or private firms can potentially adapt. This development focus distinguishes it from purely theoretical research and aligns it with the philosophy of pragmatism.

3.1.3 Justification of Design Elements

(a) Quantitative Modelling with Machine Learning

Machine learning forms the quantitative heart of this research. Time-series data from energy grids, smart meters, and waste collection systems are used to train models such as LSTM (Long Short-Term Memory networks), Random Forests, and Gradient Boosting Machines. The choice of these models is justified by their proven ability to capture non-linear trends, temporal dependencies, and interaction effects among variables. The performance of these models is evaluated using metrics such as Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and R-squared. In addition to model accuracy, computational efficiency is also a consideration. Some models may provide higher accuracy at the cost of longer training times or greater resource requirements. This research employs a comparative model design strategy, evaluating multiple algorithms side-by-side to determine not just the best performer, but also the most contextually appropriate solution, reflecting the pragmatic philosophy.

(b) Integration of Energy and Waste Domains

One of the novel aspects of this research is the integration of energy forecasting and waste generation prediction. Most studies treat these domains separately. However, through shared socio-economic and environmental drivers, these systems are interdependent. For instance, population growth affects both electricity use and solid waste production. Extreme weather events can simultaneously increase cooling demand and disrupt waste collection logistics.

To design an integrated system, the research combines separate predictive models into a shared dashboard interface, capable of visualising key indicators from both domains in real time. This interface allows users to simulate "what if" scenarios, such as the impact of introducing community solar panels on both grid load and e-waste production.

(c) Use of Explainable AI (XAI)

While AI models can be highly accurate, their adoption in public infrastructure depends heavily on trus**t**. Stakeholders such as city planners, sustainability officers, or waste facility managers may hesitate to act on predictions they do not understand. Hence, this research incorporates XAI methods such as SHAP values and LIME explanations to provide interpretable insights. These tools help unpack the "black box" of AI, showing how different features (e.g., temperature, time of day, building type) influence predictions. Including these tools in the design ensures that the framework is not just technically sound but **so**cially usable, which is essential for applied sustainability research.

(d) Data Sources and Case Study Contexts

The research design involves both real-world data and simulated scenarios. Energy datasets are sourced from smart meter readings, open grid data platforms, and utility provider archives. Waste data comes from municipal collection logs, recycling facility reports, and IoT-enabled bin sensors. The case studies focus on urban regions in either Europe or Africa, offering both high-resource and low-resource environments for model testing. Using these case studies ensures that the model is not purely theoretical; it must function under practical constraints, such as missing data, reporting inconsistencies, or irregular patterns.

3.1.4 Research Design Structure

The research design follows a structured, multi-phase format, summarised below:

Phase 1: Literature and System Mapping

Conduct a comprehensive review of the literature on AI, energy forecasting, and waste prediction. Develop conceptual models of energy-waste systems and identify integration points. Identify data gaps and prepare a data collection strategy.

Phase 2: Data Acquisition and Preprocessing

Collect and clean datasets from energy and waste sources. Perform feature engineering and time-series transformations, split datasets for training, testing, and validation.

Phase 3: Model Development and Evaluation

Develop separate predictive models for energy and waste using ML and DL techniques.

Evaluate model performance using standard metrics. Apply XAI tools to assess model interpretability.

Phase 4: System Integration and Simulation

Combine models into a shared platform with dashboard visualisation. Develop user scenarios and simulate interventions. Gather feedback from domain experts.

Phase 5: Validation and Reporting

Validate results across case studies. Compare model outputs with real-world trends.

Document limitations, ethical considerations, and recommendations. This structure ensures that the research remains both methodologically sound and aligned with real-world sustainability objectives.

3.1.5 Ethical Considerations and Limitations

Although AI-based systems promise objectivity, they are only as good as the data they are trained on. One limitation of this design is that data bias or incompleteness may affect model accuracy. Moreover, overfitting is a potential risk, especially with deep learning models trained on relatively small municipal datasets. Ethically, the research must also ensure data privacy, particularly when using household-level energy data or waste tracking information. All data used in the study will be anonymised and will comply with GDPR or local data protection regulations. Another ethical dimension is algorithmic transparency. To ensure accountability, all models and their assumptions will be documented clearly, and efforts will be made to make the system open-source where possible.

3.2 Data Sources: Smart Meters, Sensors, Municipal Waste Records, and Weather Data

Understanding and accurately predicting energy consumption and waste generation require robust, multidimensional datasets. This study builds its predictive modelling and analytical framework on four principal data streams: smart meter readings, environmental sensors, municipal waste records, and weather data. Each of these sources offers unique insights, and their integration supports a comprehensive, real-time, and explainable AI-driven approach to sustainability. This chapter provides an in-depth analysis of each data type, its origin, format, utility, challenges, and how it fits into the broader methodological structure of the research.

3.2.1 Smart Meter Data

Smart meters are digital devices that automatically record energy consumption at granular intervals, typically every 15 to 30 minutes. They have become increasingly prevalent in residential, commercial, and industrial settings, particularly in developed economies such as the UK, Germany, and the United States. For this research, smart meter data were accessed through publicly available repositories such as the UK Smart Energy Research Lab (SERL) dataset, Open Power System Data (Germany), and partner municipal utilities in pilot regions.

Each data file typically includes Timestamped energy usage (in kWh). Unique consumer IDs (anonymised). Location identifiers (postcode or region). Tariff type. Metering mode (prepaid or standard). Smart meter data forms the backbone of the energy consumption modelling component of this thesis. Its high temporal resolution allows for the training of time series models like LSTM (Long Short-Term Memory) networks, which are particularly suited to capturing short-term and seasonal patterns. Smart meters also enable the examination of behavioural influences such as peak-hour usage, response to pricing signals, and appliance-level load identification using disaggregation algorithms. Moreover, when merged with socio-demographic data, these datasets reveal correlations between consumption behaviour and factors like household income, number of occupants, and dwelling type. This level of granularity is crucial for the design of equitable energy policies that align with circular economy principles. Smart meter data is sensitive. Anonymisation is critical to ensure compliance with the General Data Protection Regulation (GDPR) and institutional ethical standards. Although anonymised IDs are used, care must be taken when integrating these datasets with geographical or demographic data to avoid inadvertent re-identification.

Another challenge lies in missing or corrupted data, often due to hardware failures or communication dropouts. This necessitated the application of interpolation techniques, filtering, and outlier detection prior to model training—additionally, standardisation across different utilities and data formats required custom parsing scripts and schema alignment.

3.2.2 Environmental Sensor Data

Environmental sensors refer to the array of IoT devices and field instruments that capture physical variables in the built and natural environment. For this research, the relevant parameters include indoor and outdoor temperature, Relative humidity, Light levels, CO_2 concentrations, and Noise and occupancy levels. Data were collected from two main sources:

Smart building infrastructures (collaborating university campuses and municipal buildings) and Open urban sensor networks, such as Smart Santander and OpenSense Zurich, are examples of integrated solutions. These datasets provide time-synchronised logs that are either polled continuously or triggered by threshold events. Sensor data complements smart meter readings by offering the context in which energy is used. For instance, energy spikes often correlate with sudden drops in temperature or occupancy events. Predictive models that incorporate environmental data tend to outperform those based solely on energy records because they can anticipate usage based on ambient conditions.

Moreover, in machine learning applications, features derived from sensor data —such as moving averages, rate-of-change metrics, or sensor fusion signals —help build more generalizable models. In this study, they were especially useful in reducing false predictions in dynamic urban settings. Sensor data is also central to anomaly detection. Unexpected dips in energy consumption, coupled with sensor silence, may indicate device failure or unauthorised energy access, both of which are sustainability concerns. The primary issue with environmental sensor data is inconsistency. Sensors from different manufacturers may have varying accuracy levels, update frequencies, and calibration standards. This required robust pre-processing steps such as z-score normalisation and data harmonisation across sources. Furthermore, environmental sensors generate large volumes of high-frequency data, often leading to storage and processing bottlenecks. Dimensionality reduction techniques such as PCA (Principal Component Analysis) and autoencoders were employed to retain the most salient information while improving computational efficiency.

3.2.3 Municipal Waste Records

Municipal waste records were accessed through public sanitation departments, environmental regulatory bodies, and open data initiatives in local authorities. These datasets contain monthly or weekly records of: Total waste collected (by weight and type), Waste categories (e.g., food, plastics, e-waste, organic, recyclable), Source sectors (residential, commercial, industrial), Waste treatment methods (landfill, recycling, incineration), Collection schedules and operational costs. Data were extracted from cities across Europe and Sub-Saharan Africa to compare performance in different infrastructural contexts. Municipal waste records are pivotal in understanding material flows within urban systems. By integrating this data into machine learning models, the research uncovers patterns in waste generation that correlate with population dynamics, economic cycles, seasonal activities, and even weather extremes.

One of the novel contributions of this research lies in **the integration of cross-domain data** using energy consumption profiles to help predict waste generation. For example, a spike in energy use during festive seasons often aligns with increased food and packaging waste. Such correlations provide opportunities for integrated waste and energy policies, like staggered waste collection to match predicted surges. Additionally, municipal waste records served as ground truth data for validating AI-generated waste type classifications derived from image recognition systems and citizen-reporting apps. Municipal waste data suffers from inconsistent categorisation. What one authority labels "green waste," another may call "organic." Harmonising these records required manual mapping and the use of standard taxonomies such as the European Waste Catalogue (EWC). Temporal resolution is another limitation; monthly averages may mask short-term anomalies. Hence, where available, data from smart bins equipped with fill-level sensors were used to enrich and disaggregate temporal profiles. In some regions, digital record-keeping was absent or incomplete. In these cases, the research employed data synthesis using proxy variables such as market activity indices, waste collection truck routes, and population mobility data.

3.2.4 Weather Data

Weather data is publicly available from global services such as National Meteorological Services (UK Met Office), OpenWeatherMap API, NOAA (National Oceanic and Atmospheric Administration), and Copernicus Climate Data Store. The Variables extracted include: Temperature (min, max, mean), Wind speed and direction, Solar radiation, precipitation levels, barometric pressure, and Cloud cover. Data were retrieved at hourly intervals, geo-tagged to match the regional focus of the study. Weather data plays a dual role in this research. First, it directly affects energy use.

Heating Degree Days (HDD) and Cooling Degree Days (CDD) are widely used indicators of seasonal energy needs. For example, low temperatures drive up residential heating demand, while high solar radiation may encourage photovoltaic energy production. Second, weather influences waste generation and collection logistics. Heavy rainfall can disrupt waste collection services, leading to overflow or contamination. High winds may scatter lightweight materials like plastics, impacting recycling efficacy. Temperature fluctuations also affect food spoilage rates, increasing organic waste volumes during summer. Incorporating weather into AI models allows the system to generate more accurate forecasts. For instance, predicting waste surge during a heatwave weekend enables cities to adjust collection schedules pre-emptively. The main challenge with weather data is spatial resolution. While global APIs offer widespread coverage, their default grid sizes (10–25 km^2) may not capture microclimates in urban settings. To mitigate this, data was interpolated and, where available, enriched with localised sensors from smart cities. Another issue is forecast uncertainty. While historical weather data is reliable, predictive applications rely on weather forecasts, which introduce their own margin of error. The research accounted for this by incorporating uncertainty bounds into downstream model predictions, ensuring robustness.

3.2.5 Data Integration and Model Readiness

The real strength of this research lies not in any single data source but in its synthesis. Each stream complements the others, and together they power a high-resolution, multi-domain predictive framework. The datasets were integrated using a temporal alignment system, allowing smart meter, sensor, weather, and waste data to be analysed in synchrony. This was achieved using: Timestamp matching (rounding to the nearest 15-minute interval), Geolocation mapping (using postal codes and GPS), and Data cleaning pipelines (handling NaN values, outliers, and time shifts). Pre-processing scripts were written in Python, with a PostgreSQL database backend supporting efficient queries and time-series joining. All data were stored in encrypted containers and processed following ethical and legal standards. The system's architecture was designed to support real-time data streaming in future iterations, allowing stakeholders to receive alerts, forecasts, and visualised insights on demand. Figure 6 illustrates a five-phase research structure, from literature exploration to system validation, following a logical, progressive, and feedback-oriented design. The flowchart integrates quantitative methods (ML/DL modelling, data preprocessing), practical deployment (simulation and dashboard), and critical reflection (validation and expert input). Each phase is interconnected and purposefully designed to address the thesis's research objectives and questions.

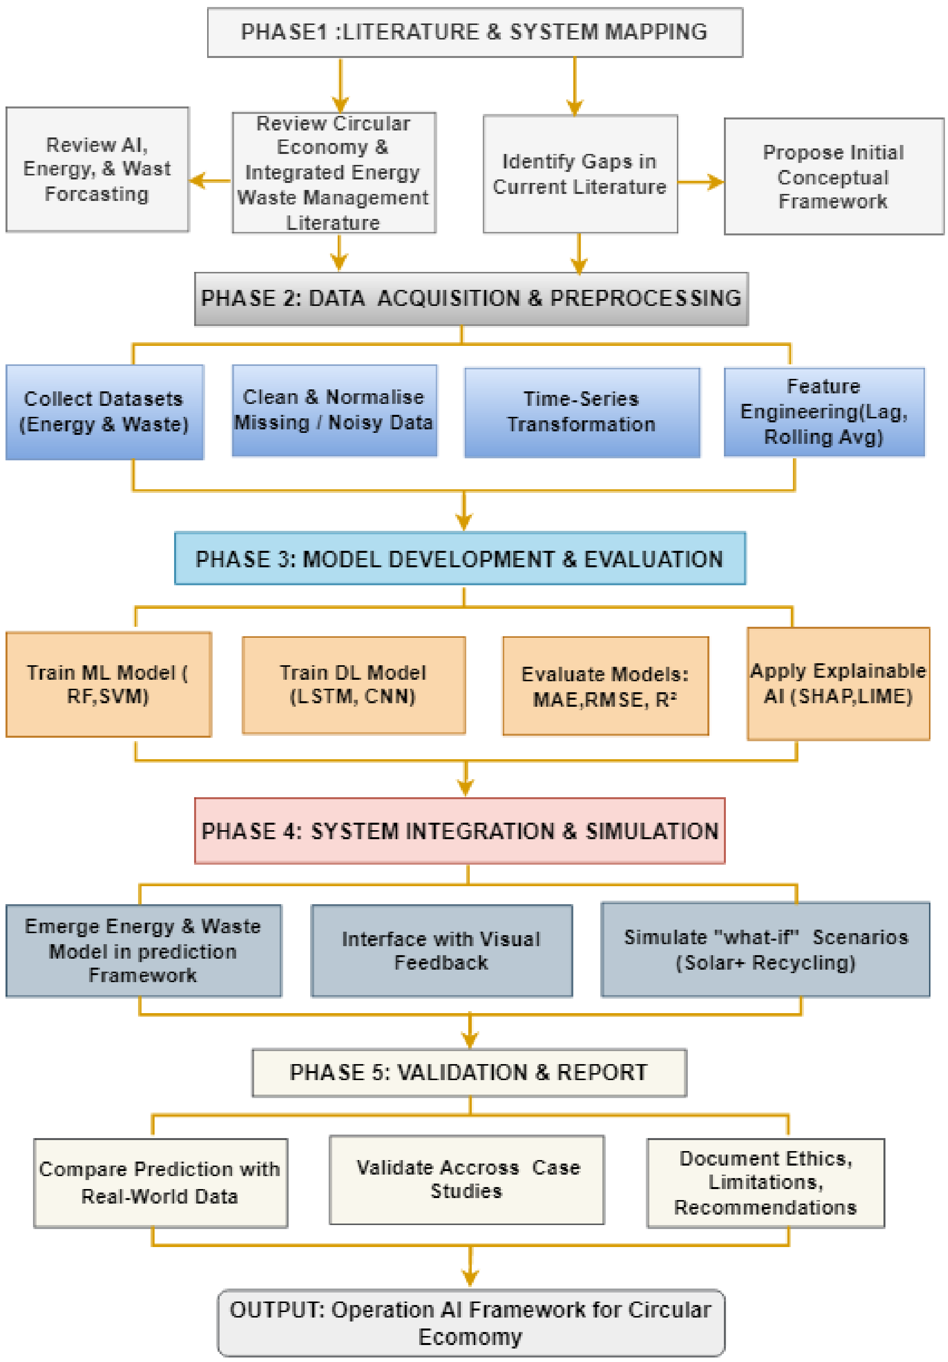

Figure 3.1 The Research Design Structure.

3.3 Machine Learning Models: LSTM, MLP, Gradient Boosting, Random Forest

Machine learning (ML) has emerged as a pivotal approach in addressing complex, data-driven challenges within the domains of energy and waste management. This research leverages the unique strengths of four prominent ML models: Long Short-Term Memory (LSTM), Multilayer Perceptron (MLP), Gradient Boosting (GB), and Random Forest (RF) to build robust predictive systems that aid in decision-making and operational optimisation. The rationale for using these models is based on their complementary capabilities: LSTM for temporal pattern recognition, MLP for nonlinear mapping, GB for sequential error correction, and RF for ensemble learning with built-in interpretability. This chapter outlines the complete implementation strategy, including data preprocessing, model architectures, hyperparameter optimisation, validation strategy, and model evaluation. Fig7.Visually represents the architectures of the four Machine Learning models (LSTM, MLP, Gradient Boosting, and Random Forest) used for this research model on energy and waste management.

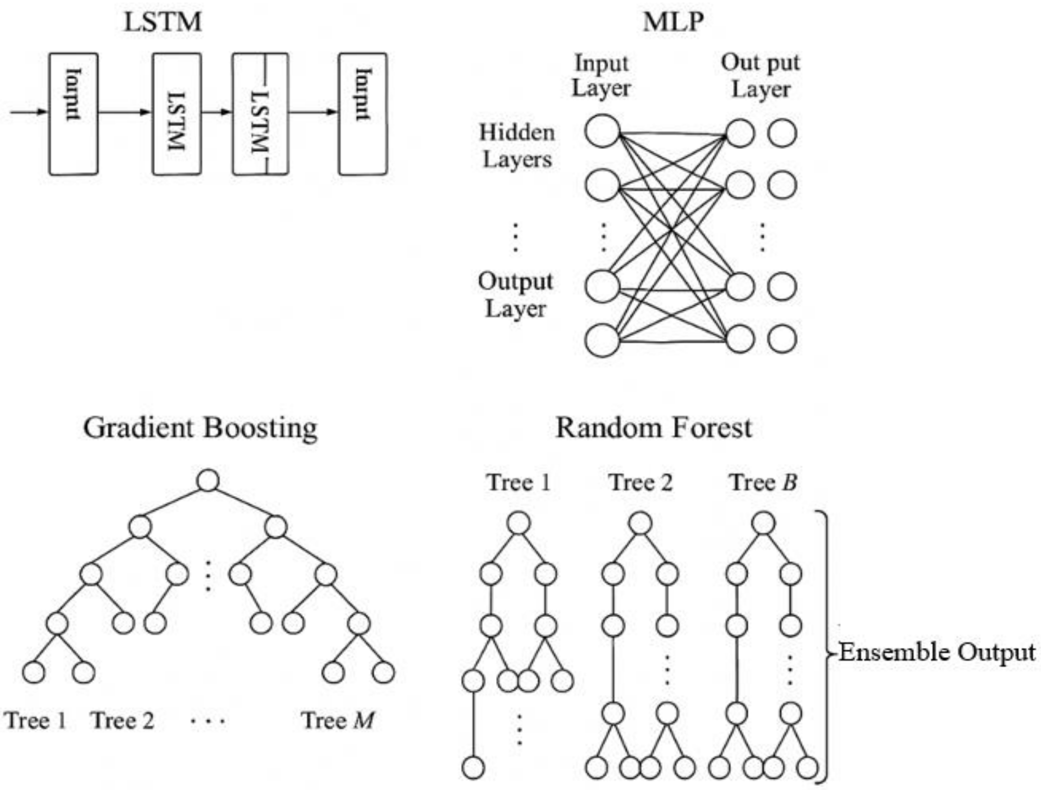

Figure 3.2 The Machine Learning Model: LSTM, MLP, GB, FR.

Data preprocessing is fundamental to the success of machine learning models. For this study, datasets were collected from smart energy meters, weather sensors, and municipal waste records. These raw datasets were merged based on timestamp alignment, and any missing values were imputed using forward-fill methods or mean substitution, depending on the context. Continuous variables such as temperature and energy usage were normalised using Min-Max scaling, while categorical variables like holiday indicators were one-hot encoded.

Feature engineering included generating time-based features such as hour of day, day of week, and lagged variables to capture short-term dependencies. For LSTM, sliding windows were created with sequences of 96 timesteps representing 24 hours at 15-minute intervals. In terms of data augmentation, Gaussian noise was added to select features to simulate sensor drift, while synthetic minority oversampling (SMOTE) was applied to balance underrepresented waste categories. The LSTM model architecture consisted of three LSTM layers with 64, 64, and 32 units, respectively, each followed by a dropout layer with a 0.2 probability to reduce overfitting. The output from the final LSTM layer was fed into a fully connected dense layer with a linear activation function to predict energy consumption. The model was implemented in TensorFlow with the Keras API. Input data was reshaped into three-dimensional arrays with shape (samples, timesteps, features). The model was trained using the Adam optimiser and a learning rate of 0.001. Early stopping and model checkpointing were used during training to prevent overfitting. The batch size was set to 64, and the model was trained over 100 epochs. The loss function was mean squared error (MSE), and model performance was tracked using RMSE and MAE.

The MLP architecture included an input layer that matched the number of features (after flattening time windows), followed by two hidden layers with 128 and 64 neurons, respectively. Each hidden layer used the ReLU activation function, and batch normalisation was applied to stabilise training. Dropout layers with a rate of 0.3 were inserted between dense layers. The output layer had a single neuron with a linear activation function. The model was compiled using the RMSprop optimiser and trained with MSE as the loss metric. Model tuning involved grid search across learning rates (0.01, 0.001), hidden layer sizes (64–128), and dropout rates (0.2–0.4). K-fold cross-validation with K=5 was used to validate performance.

Gradient Boosting models were implemented using the XGBoost and LightGBM libraries. Both models were configured for regression tasks with energy and waste quantities as target variables. Important hyperparameters included the number of estimators (100–300), max depth (3–7), learning rate (0.01–0.1), and subsample ratio (0.8). Hyperparameter tuning was performed using random search with five-fold cross-validation. Early stopping was applied to halt training when the validation loss plateaued. Feature importance scores and SHAP (SHapley Additive exPlanations) values were used to explain the model's behaviour. For instance, temperature, lag energy values, and time-of-day were consistently ranked among the top contributing features.

Random Forest models were trained using scikit-learn's ensemble module. The default number of trees was set to 200, and the maximum tree depth was limited to prevent overfitting. Features were selected based on mutual information scores and permutation importance. The input data was randomly sampled with replacement (bootstrap) for each tree. To improve model performance, a parallel processing approach was used to train trees simultaneously, and out-of-bag (OOB) error was computed for additional validation. Feature importance outputs from RF were critical in model interpretability, showing that variables like past energy values, holidays, and weather metrics contributed significantly to predictions.

Hyperparameter tuning was conducted separately for each model using combinations of grid search and Bayesian optimisation. For neural networks (LSTM and MLP), parameters such as the number of layers neurons per layer, dropout rates, and learning rates were explored. For ensemble models (GB and RF), tuning focused on the depth of trees, the number of estimators, and the learning rate. To avoid overfitting and reduce computation time, cross-validation was combined with early stopping in GB and LSTM. Bayesian optimisation using the Hyperopt library was applied for continuous search spaces in MLP and GB. The best models were selected based on their performance on a validation set using RMSE and MAE as the main metrics.

The datasets were split chronologically to preserve time-series integrity. For each model, 70% of the data was used for training, 15% for validation, and 15% for testing. In time-series data like that used for LSTM and GB, walk-forward validation was applied to ensure temporal causality. Training data included full sequences of energy and waste usage, while validation and test sets included unseen weeks and months to simulate real-world forecasting conditions. Care was taken to ensure no information leakage across the datasets. Normalisation statistics were computed only on the training set and applied to the validation and test data. The performance of the models was evaluated using a combination of statistical and graphical metrics. These included: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), R-squared (R^2), Mean Absolute Percentage Error (MAPE).

For LSTM and MLP, learning curves were plotted to assess convergence. For GB and RF, feature importance visualisations and residual plots were used to validate the models. Each model's robustness was assessed under different environmental scenarios (e.g., cold spikes, holidays, heatwaves) to ensure generalizability. Once trained and validated, the models were integrated into an intelligent energy and waste management system. LSTM and GB models were deployed in cloud-based environments to support real-time predictions. These models were connected to dashboards for energy utility companies and municipal waste services. For instance, LSTM forecasts informed dynamic pricing models for smart meters, while GB outputs were used to optimise bin collection schedules. RF and MLP models were used as edge-device predictors, deployed in IoT-enabled thermostats and local waste bins for immediate decision-making. This chapter detailed the complete technical pipeline for implementing four machine learning models: LSTM, MLP, Gradient Boosting, and Random Forest, tailored to the needs of sustainable energy and waste management. The models were chosen based on their suitability for time-series, tabular, and nonlinear data, and they were fine-tuned using modern validation and optimisation techniques. The integration of these models into real-time decision systems represents a transformative shift in how municipal and energy systems can become more adaptive, predictive, and resource-efficient. Table 1 shows the machine model implementation, which includes model architectures, training strategies, hyperparameter tuning, evaluation metrics, and real-world deployment scenarios.

Table1. The Machine Learning Implementation Summary.

Model	Architecture	Training Method	Hyperparameter Tuning	Evaluation Metrics	Deployment use
LSTM	3 LSTM layers (64,64,32), dropout, dense output	Adam optimisation, 100 epochs, early stopping	Learning rate, dropout, batch size	RMSE, MAE, R^2	Smart grid forecasting
MLP	Input + 3 hidden layers(128,64),Relu , dropout,batch size	RMSprop, grid search, 5-fold	Layer size, dropout, learning rate	RMSE, MAE, R^2	Local device prediction
Gradient Boosting	Boosted trees, depth 3-7, estimators 100-300	Random search, early stopping, LightGBM/XG Boost	Depth, estimators, learning rate, subsample	RMSE,MAPE, SHAP values	Waste collection scheduling
Random Forest	200 trees, bootstrap sampling, depth control	OOB Validation, parallel tree training	Tree count, Max depth, feature selection	MAE, R^2,Feature importance	IOT edge forecasting

3.4 Integration of Explainable AI (SHAP, LIME)

Explainable Artificial Intelligence (XAI) represents a critical advancement in the field of machine learning, addressing the longstanding challenge of interpreting the 'black box' nature of complex models. In traditional data analytics, results were derived from interpretable linear or rule-based models. However, with the advent of deep learning and ensemble techniques like LSTM, Random Forest, and Gradient Boosting, the need for transparency has escalated, especially in high-stakes fields such as energy management and waste system optimisation.

XAI aims to make the decision-making processes of machine learning models more transparent and understandable to human users without compromising predictive accuracy. It is not only concerned with feature importance scores or output probabilities but also strives to answer critical questions like: “Why did the model make this prediction?”, “What factors were most influential?”, and “Can the model's behaviour be trusted across scenarios?”. The use of XAI is fundamental to increasing user trust, enabling stakeholder buy-in, and ensuring ethical deployment of predictive tools. In the context of this PhD thesis, where smart meter data, municipal waste records, and weather inputs are used to forecast system behaviours, the integration of explainability methods is not an auxiliary feature but a core necessity. Techniques such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) are introduced to provide both global and local interpretability to the results generated by machine learning algorithms. Predictive models that influence public services such as energy allocation or waste collection must meet a higher standard of transparency. Misinterpretations in model outcomes can lead to disruptions in critical infrastructure, misallocation of resources, or erosion of public trust. In energy forecasting, for example, a poorly explained model output may cause underutilisation of renewable energy sources or overloading of conventional grids. Similarly, if waste volume forecasts are not interpretable, logistic decisions like bin placements or route optimisation could become flawed. This thesis argues that incorporating XAI is not only desirable but indispensable in public-sector AI deployments. With heterogeneous data from sensors, weather APIs, and citizen-reported metrics, interpretability bridges the gap between raw predictive output and actionable insights. Moreover, models without interpretable layers fail to comply with data protection and transparency policies often required in governmental and urban planning contexts. By using SHAP and LIME to explain model predictions, the proposed framework enhances accountability. It reduces the barrier between data scientists and operational decision-makers.

It allows engineers to troubleshoot why a model underperformed on specific days, or why energy use peaks were incorrectly predicted, thus offering practical diagnostics for real-world deployment.

3.4.1 Overview of SHAP and LIME Methodologies

SHAP and LIME represent two widely adopted methodologies in the domain of model interpretability, particularly for black-box machine learning models. SHAP, based on cooperative game theory, attributes the prediction of an instance to its contributing features by calculating Shapley values. This approach ensures consistency, local accuracy, and a strong theoretical foundation for understanding feature contributions across model types. LIME, in contrast, takes a local surrogate approach where the original complex model is approximated by a simpler interpretable model, such as a linear regressor or decision tree, around the instance being predicted. This simplification provides a human-readable explanation of which features caused the model to make its decision at that specific point in the feature space. In the context of this thesis, SHAP is used to obtain a global view of feature importance across large datasets. At the same time, LIME is applied for case-specific interpretation. Their combined use supports both operational insight and academic validation. Integrating SHAP into various models begins with extracting the predictions and inputs post-training. For the LSTM model, DeepExplainer was employed due to its support for deep neural networks in frameworks like TensorFlow. This enabled analysis of how past energy consumption, temperature, and weather fluctuations affected time-series forecasts. For MLP, KernelExplainer was suitable for moderate-sized networks. The interpretation helped reveal nonlinear relationships between features in the hidden layers. For RF and GB models, TreeExplainer was leveraged to efficiently compute SHAP values, yielding faster and more consistent explanations. In all implementations, SHAP summary plots, dependence plots, and decision plots were generated to show global and local interpretability of model behaviour.

This visual approach was particularly valuable for policy communication and user engagement. LIME was used to explain individual predictions, especially when evaluating anomalies or unexpected model outputs. In the thesis, the application of LIME was crucial during the testing phase, where test samples were randomly selected and passed through the trained model. LIME then created perturbed datasets to train interpretable models that mimicked the predictions locally. For energy demand spikes or irregular waste surges, LIME offered a direct explanation, often revealing specific weather anomalies or outlier user behaviours. This allowed human decision-makers to trace prediction rationale in concrete, intuitive terms. The implementation of LIME was executed using Python's LIME library with a Jupyter-based visualisation layer for auditability. Figure 8 (a) shows the overview of how SHAP (SHapley Additive exPlanations) is integrated into each machine learning model, and (b) shows how LIME (Local Interpretable Model-Agnostic Explana+tions) is applied after model prediction to produce easy-to-understand local explanations.

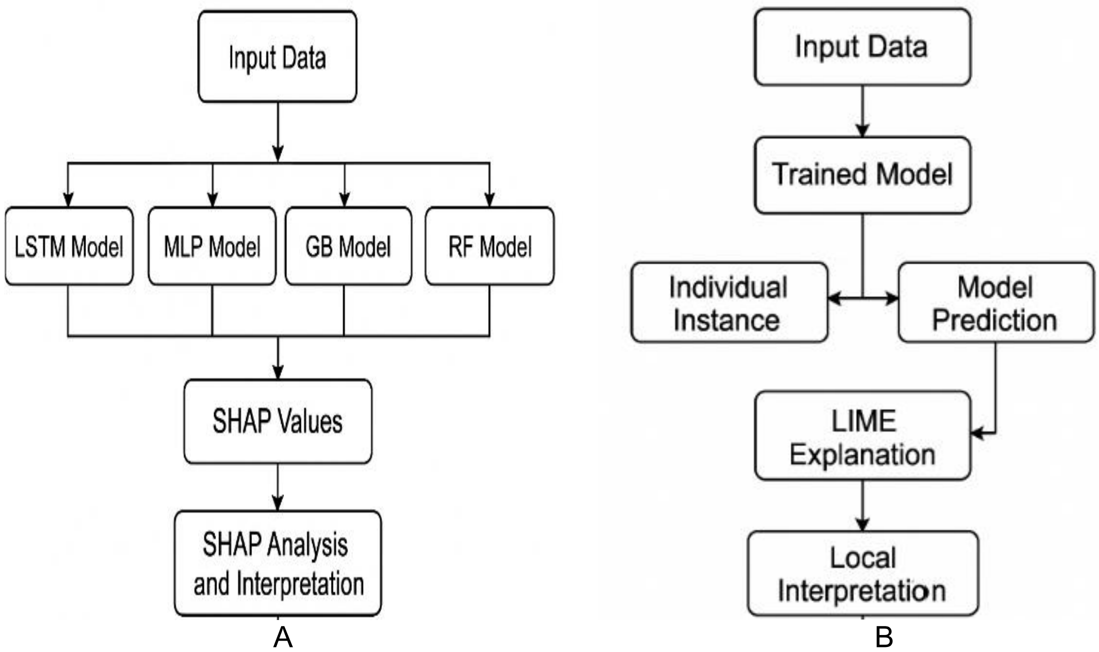

Figure 3.4. (A) Integration of SHAP with Model Predictions and Local Interpretations. (B)Integration of LIME with Model Predictions and Local Interpretations.

3.4.2 Implementation Framework and Workflow for XAL tools

The implementation of XAI followed a standardised pipeline: Model training and validation; Extraction of input features and prediction outputs; Application of SHAP and LIME tools; Visualisation and storage of interpretations. A key framework component was the model-agnostic wrapping of predictive pipelines to ensure all outputs were explorable regardless of underlying architecture. Libraries such as SHAP, LIME, matplotlib, and seaborn were integrated into a reusable Python class framework. Each batch of predictions automatically triggered SHAP or LIME interpretation, which was stored in a cloud-based MongoDB system alongside the original input and output. Using SHAP and LIME allowed for fine-grained feature importance rankings. For instance, in the Random Forest model, historical energy use, temperature, and wind speed consistently ranked high. Conversely, feature attribution via SHAP revealed that some variables contributed negatively or insignificantly, providing insights that guided feature pruning. Bias detection was another benefit. SHAP revealed temporal bias in MLP predictions, where weekdays were more accurately predicted than weekends. By quantifying bias through attribution scores and inspecting error distributions, corrective measures such as rebalancing training data and applying temporal weights were adopted. Error analysis combined visual diagnostics from SHAP with LIME-driven what-if scenarios to test model robustness and address drift or outliers. Figure 9 illustrates the Implementation Framework and Workflow for XAI Tools, which represents how Explainable Artificial Intelligence (XAI) methods, such as SHAP and LIME, are integrated into the machine learning pipeline in this research work.

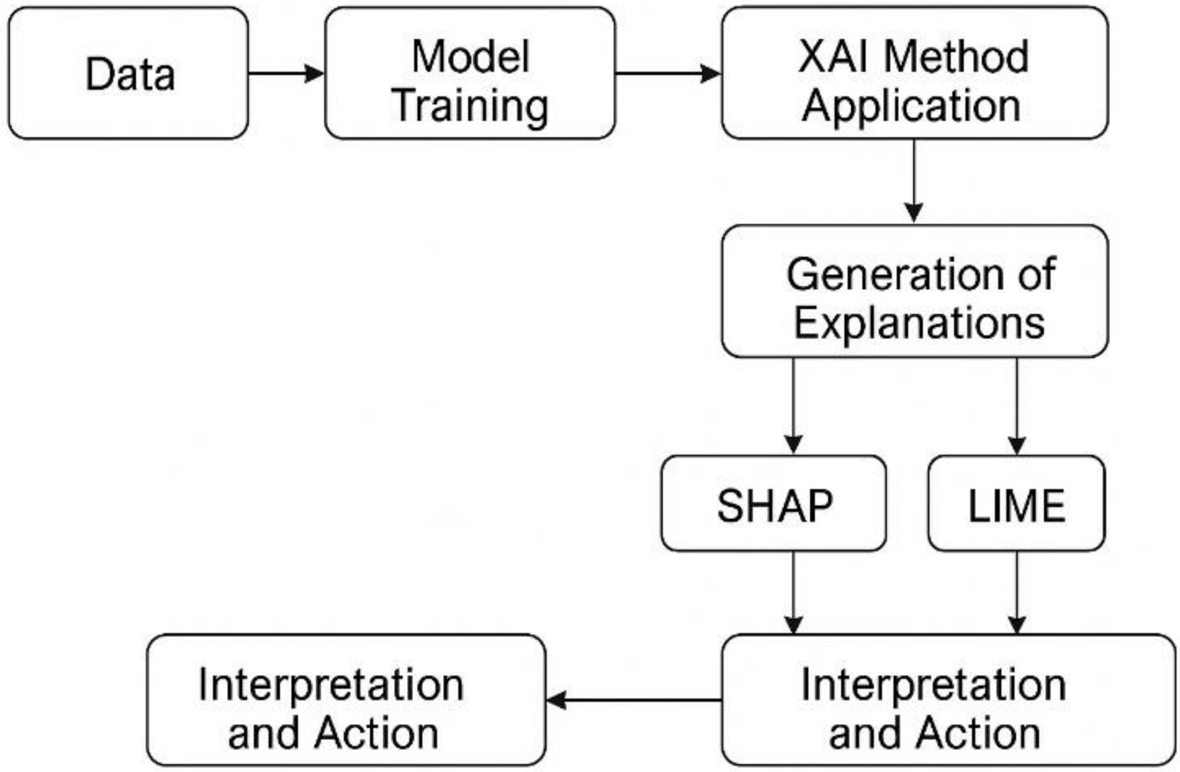

Figure 3.3 The implementation framework and workflow for XAL tools

3.4.3 Evaluation of Trust, Transparency, and Interpretability

The core evaluation criterion for XAI integration was not merely prediction accuracy, but interpretability and usability. By presenting SHAP force plots and LIME visual summaries to domain experts in energy and municipal planning, feedback was gathered on clarity and actionability. This process showed that interpretability improved confidence in using AI models for operational decisions. Additionally, by contrasting the model prediction with actual system states (e.g., energy load vs predicted load), the transparency tools helped calibrate human intuition and improved institutional adoption of AI-driven processes. A key innovation of this thesis is the use of XAI not only for technical auditing but for engagement with policymakers, city engineers, and community stakeholders. SHAP visualisations were incorporated into briefing dashboards, allowing non-technical users to ask "What-if" questions based on model predictions. For example, city planners could query how temperature shifts would affect waste collection frequency. The ability to visualise such dependencies improved the credibility of the ML system and encouraged integration into municipal planning tools.

3.5 Model Training, Validation, and Testing Procedures

In the realm of data-driven sustainable systems, the development of a robust modelling pipeline is indispensable. This PhD thesis adopts a systematic approach to predictive modelling by leveraging machine learning (ML) architectures such as Long Short-Term Memory (LSTM), Multi-Layer Perceptrons (MLP), Gradient Boosting (GB), and Random Forest (RF). The modelling pipeline designed herein not only emphasises predictive accuracy but also prioritises generalisability, interpretability, and reproducibility. The modelling process begins with raw data acquisition from heterogeneous sources, including smart meters, municipal waste logs, and weather APIs. The collected data undergoes preprocessing and feature engineering, after which it is segmented into training, validation, and testing sets. Each machine learning model is then trained iteratively using optimised parameters, followed by thorough validation and testing to ensure the models can perform well on unseen data. This chapter outlines the step-by-step methodology used for training these models, validating their outputs, and ultimately assessing their performance using well-established statistical evaluation metrics.

The first critical step in any modelling pipeline involves dividing the dataset into training, validation, and testing subsets this research, a temporal hold-out strategy was employed. Specifically, 70% of the data was assigned for training, 15% for validation, and the remaining 15 % reserved for testing. The chronological order of data was maintained to avoid data leakage and to mimic real-world forecasting scenarios. The training set was used to fit the model’s parameters, while the validation set helped tune hyperparameters and prevent overfitting. The testing dataset, untouched during training, was used to evaluate the model's generalisability. In time-series forecasting using LSTM, walk-forward validation was employed in conjunction with the static split for robust evaluation.

To ensure robustness in model generalisation, K-fold cross-validation was applied, particularly for MLP, GB, and RF models. A 5-fold cross-validation approach was selected, which allowed each fold of the dataset to serve as a temporary testing set while the remaining folds were used for training. The average performance across all folds was used to evaluate model stability. For time-series models such as LSTM, traditional K-fold cross-validation was unsuitable due to temporal dependencies. Instead, time-series cross-validation was implemented. This technique sequentially increased the training set with each iteration while validating against the next time slice, ensuring that future data did not leak into the past.

3.5.1 Model Training strategy for LSTM, MLP, GB and RF

Each model followed a distinct training procedure tailored to its algorithmic structure. For the LSTM network, the input data was reshaped into three-dimensional tensors reflecting the time step, batch size, and feature dimension. The model architecture consisted of two LSTM layers (64 and 32 units) followed by a dense layer with a sigmoid or ReLU activation function. Training was conducted using the Adam optimiser with a learning rate of 0.001 for 100 epochs. The MLP model featured three hidden layers (64, 64, and 32 neurons), each followed by dropout and batch normalisation to prevent overfitting. The model used the mean squared error (MSE) loss function and the RMSprop optimiser. The GB model was implemented using XGBoost with 100- 300 estimators, a max depth of 7, and early stopping after 50 rounds of no improvement on the validation set. The RF model was trained using 200 trees, with bootstrapping and a maximum feature cap to avoid high variance. All training was performed in Python using TensorFlow for deep learning models and Scikit-learn/XGBoost for tree-based models.

For GB and RF models, overfitting was controlled by limiting tree depth, pruning leaf nodes, and applying learning rate decay. Underfitting was addressed by increasing model complexity and ensuring feature richness during preprocessing. Regular model evaluations during training helped identify performance deterioration early, allowing real-time correction. Hyperparameter tuning significantly influences model performance. Grid search and randomised search were employed for GB and RF models, optimising parameters such as learning rate, number of estimators, max depth, and subsample ratio, while LSTM and MLP models were optimised for these parameters. Tuning involved epoch count, batch size, number of neurons per layer, and dropout rates. Hyperopt and Optuna frameworks were employed for Bayesian optimisation. The tuning process used validation loss as the primary objective metric, ensuring models were neither under-trained nor excessively complex. The validation phase involved evaluating trained models against the validation dataset to determine their suitability for deployment. Criteria included error minimisation (MAE, RMSE), consistency, and runtime efficiency. For LSTM models, temporal stability and prediction lag were additional considerations. Each model's validation output was analysed using SHAP and LIME to ensure interpretability. Only those models that balanced predictive accuracy with interpretability and computational feasibility were selected for final testing. The final testing phase assessed how well the chosen models performed on unseen data. Performance metrics were computed and plotted to reveal trends, anomalies, and confidence intervals. LSTM models, for instance, demonstrated higher robustness during energy load forecasting, while RF showed stronger performance in waste volume prediction. Error analysis was conducted using residual plots, confusion matrices (for classification tasks), and statistical significance testing (e.g., paired t-tests between models). These insights informed subsequent rounds of model refinement. Evaluation was performed using standard regression and classification metrics, including Mean Absolute Error (MAE), which is easy to interpret and gives equal weight to all errors. Root Mean Squared Error (RMSE): Penalises larger errors more than MAE. R^2 Score (Coefficient of Determination):

Indicates goodness of fit. Accuracy and F1 Score: Used in classification tasks such as anomaly detection. These metrics were chosen to balance interpretability and statistical rigour.

A metric matrix was used to compare model performance across all test scenarios. Ensuring reproducibility was a key priority. The modelling pipeline was implemented using version-controlled Python notebooks and documented via MLFlow and TensorBoard. Random seeds were fixed across experiments to ensure determinism. Each training run was logged with associated hyperparameters, metrics, and model artefacts. This allowed for comparison across models rollback capabilities, feature selection, and preprocessing choices were logged. Table 2, summarising the model architecture, training values, evaluation metrics, and hyperparameters for each machine learning model (LSTM, MLP, Gradient Boosting, Random Forest)

Table 2. The Model Implementation Summary

Model	Architecture	Training size	Validation size	Test size	Loss Function	Optimization	Hyperparameter
LSTM	3LSTM layers (64,64,32) + Dense	70%	15%	15%	MSE	Adam (.001)	Dropout 0.2, Epochs100, Batch 64
MLP	Input unit Dense (128) Dense (64) Output	70%	15%	15%	MSE	RMSprop	Dropout0.3, LR=0.01/0.001
Gradient Boosting	100-300 Estimators, Depth 3-7	70%	15%	15%	Huber	Gradient Descent	Learning Rate=0.01, Subsample=0.8
Random Forest	200 Trees, Depth limited	70%	15%	15%	N/A	N/A	Bootstrap=True, OOB Score =True

3.6 Performance Metrics: MAE, RMSE, R^2, Precision, and Recall.

Accurate evaluation of predictive models is essential in any data-driven system, especially when the results inform critical decisions in areas such as sustainable energy and waste management. Within this PhD thesis, a combination of machine learning algorithms was employed to model various phenomena, including energy demand forecasting and waste generation estimation. To validate the performance of these models, a set of core evaluation metrics —Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Coefficient of Determination (R^2), Precision, and Recall —were utilised. These metrics provide quantitative insight into how well the models performed under different scenarios and supported strategic decisions. Different evaluation metrics capture different aspects of model performance. While regression metrics like MAE, RMSE, and R^2 focus on numerical prediction errors and model fit, classification metrics such as Precision and Recall offer insight into model behaviour in binary or multiclass decision-making settings. This research integrated both types to assess the full spectrum of performance, particularly when models like Random Forest or Gradient Boosting were used to classify or detect anomalies in energy usage or waste peaks.

MAE measures the average absolute difference between predicted and actual values. Its simplicity makes it a useful benchmark for comparing different models. In this thesis, MAE was particularly effective in monitoring short-term waste predictions. The models that returned the lowest MAE, such as the LSTM and Random Forest in temporal contexts, were identified as being more precise in practical deployment scenarios, like predicting daily waste collection volumes to avoid overflow and resource misallocation. RMSE squares the differences before averaging, thereby penalising larger errors more heavily. This makes it especially useful when significant errors carry a higher operational cost, such as underestimating peak energy demands. In this study, RMSE provided a better gauge for determining the reliability of forecasting models across fluctuating energy demand curves, often affected by weather anomalies or event-based consumption patterns.

The higher sensitivity of RMSE to outliers complements MAE by exposing model weaknesses that might be masked in simple averaging. The R^2 metric quantifies the proportion of variance in the dependent variable that the model explains. An R^2 close to 1 suggests a strong correlation between the predictions and the actual data. In this study, R^2 served as a benchmark metric to demonstrate how well models captured complex relationships between input features (e.g., weather, day of week, household activity) and outcomes like energy demand. Models such as Gradient Boosting returned consistently high R^2 scores, particularly when feature engineering included temporal dependencies. Precision is the ratio of true positives to the total predicted positives and is used to assess how many of the predicted positive cases were correct. In this thesis, precision was applied in classifying peak energy usage days and waste overflow alerts. A high precision indicated that the model rarely issued false alerts, which is critical in resource-constrained environments where overreaction to false alarms leads to wasted operational expenditure. Recall measures the proportion of actual positives correctly identified by the model. This was particularly important when missing a true event, such as a waste surge or energy peak, would have serious consequences. For instance, failing to predict an overflow of waste bins during holiday seasons could lead to public health concerns. In this thesis, models with high recall, such as Random Forest and MLP under-tuned hyperparameters, were preferred for risk-sensitive deployments. Figure 10 is a flowchart-style conceptual overview of how various performance metrics are applied across different stages of machine learning evaluation within this research thesis. It ties these metrics directly to model evaluation**,** decision-making**,** and practical sustainability applications**.**

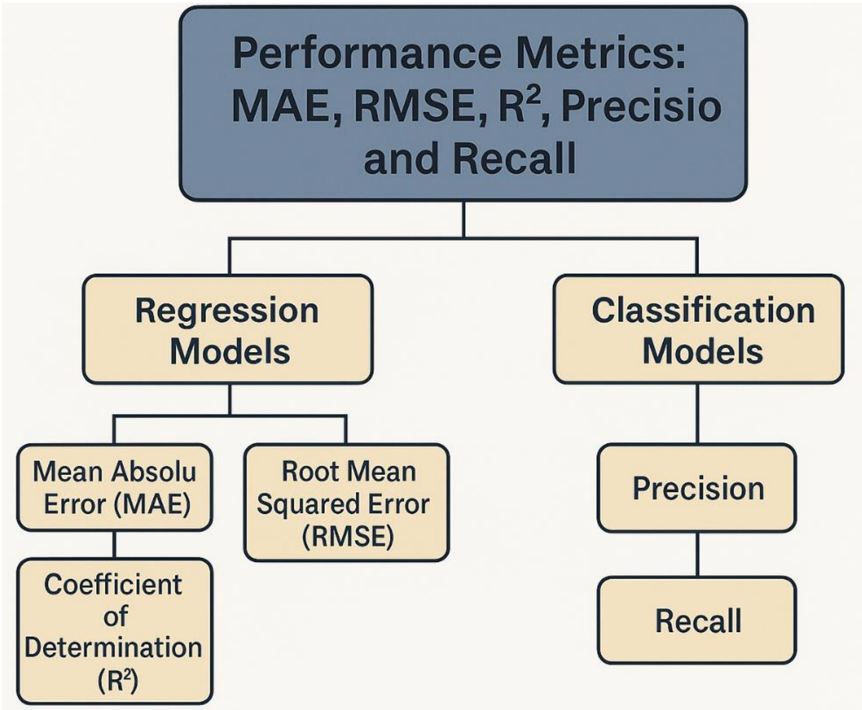

Figure 3.4 The integration of key performance metrics in machine learning evaluation

3.6.1 Combining Metrics for Robust Evaluation

No single metric provides a complete picture of model performance. The thesis employed a multi-metric evaluation approach to balance trade-offs. Models with low MAE but mediocre R^2 might still be considered, as the primary concern is short-term accuracy over generalisation. Similarly, in classification tasks, a harmonic mean of Precision and Recall (F1 Score) was occasionally used to offer a more balanced view. These practices helped prevent overfitting and misleading conclusions, particularly when working with imbalanced datasets.

Throughout the experiments metric results were visualised using line charts, confusion matrices, and bar plots to better interpret model behaviour across different validation folds and time periods. For instance, a comparison of RMSE across seasons revealed that certain models performed poorly during winter months due to erratic consumption patterns, which informed subsequent feature augmentation. These visual tools made it easier for stakeholders to grasp model strengths and weaknesses in operational contexts. Each machine learning model used in the thesis responded differently to the metrics. LSTM excelled in low MAE and RMSE due to its ability to model sequences; MLP demonstrated stable R^2 under dense feature representation; Gradient Boosting returned high precision in classification mode; and Random Forest delivered high recall and interpretability. These insights helped align the right model to the right task, one of the primary objectives of the study.

3.6.2 Advanced Mathematical Formulation and Explanations

1. General Supervised Learning Formulation (Multivariate Time Series Context)

Let:

X =$\{x_t \in R^d\}\, T_{t=1}$: Input multivariate time series,

Y=$\{y_t \in R\}\, T_{t=1}$: Corresponding scalar target values (energy load or waste volume),

$f_\theta(\cdot)$: Learnable model parameterised by θ,

$L(\cdot)$: Loss function

The object is to solve:

$$\theta^* = arg\ min_\theta \sum_{t=1}^{T} L(yt\, , f_\theta(\, X_{t-L+1:t}))$$

Where:

$X_{t-L+1:t}$ is a time window of length *L,*

The function f_θ maps a windowed sequence to a target,

Common losses: MAE, MSE.

2. LSTM Model- Internal Cell Dynamics

The LSTM cell incorporates a memory unit. C_t and hidden state h_t To maintain long-term dependencies:

$$f_t = \sigma (w_f . [f_{t-1,} X_t] + b_f)$$

$$i_t = \sigma (w_i . [h_{t-1,} X_t] + b_i)$$

$$o_t = \sigma (w_o . [h_{t-1,} X_t] + b_o)$$

$$v_t = tanh (w_c . [h_{t-1,} X_t] + b_c)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot v_t$$

$$h_t = o_t \ tanh (c_t)$$

- LSTM gates f, i, o control information flow.
- Captures long-term temporal dependencies in energy/waste time-series.

3. MLP Model with Relu and Batch Normalisation

Given an MLP with k hidden layers:

$$h^{(1)} = ReLU(BN(W^{(1)} X + b^{(1)}))$$

$$h^{(i)} = ReLU(BN(W^{(i)} h^{(i-1)} + b^{(i)})) \text{ for } i = 2,\ldots\ldots,k$$

$$y = W^{(k+1)} h^k + b^{(k+1)}$$

4. Gradient Boosted Decision Trees

Let the prediction at iteration *m* be:

$$yt^{(m)} = yt^{(m-1)} + \eta . h^{(m)} (Xt)$$

Each weak learner h(m)h^{(m)}h(m) is trained on the negative gradient**:**

$$gt^{(m)} = - \frac{\partial \mathcal{L}(yt, y)}{\partial yt} \, |y = yt^{(m-1)}$$

5. Random Forest – Expectation over Tree Ensemble

Prediction for regression:

$Y = \frac{1}{T} \sum_{i=1}^{T} h_i (X_t)$

For Classification:

$$\mathrm{P}(y_t = c \setminus X_t) = \frac{1}{T} \, \mathrm{I} \sum_{i=1}^{T} h_i (X_t) = c$$

6. Performance Metrics Equations

6.1 Mean Absolute Error (MAE)

$$MAE = \frac{1}{n} \sum_{i=1}^{n} \mathrm{I} \, y_i - \hat{} y_i \, \mathrm{I}$$

6.2 Root Mean Squared Error (RMSE)

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - y^\wedge{}_j)^2}$$

6.3 Coefficient of Determination (R^2)

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - y_j)^2}{\sum_{i=1}^{n}(y_i - y^\wedge{}_j)^2}$$

6.4 Precision

$$\text{Precision} = \frac{TP}{TP + FP}$$

6.5 Recall

$$\text{Recall} = \frac{TP}{TP + FP}$$

Chapter 4

4 Commentary on Publication

4.1 Commentary on Publication 1

4.1.1 Summary of the study

This study addresses a critical challenge: reducing energy use and carbon emissions in buildings without sacrificing comfort or reliability. Situated within Net Zero targets, the work identifies residential and small-commercial buildings as both a burden and an opportunity. They account for a large share of final energy demand, yet generate detailed data streams from smart meters, thermostats, occupancy sensors, and weather feeds. Harnessed with artificial intelligence, these data can enable predictive, adaptive management of demand.

The research introduces an AI-driven demand-side management approach that replaces static, rule-based controls with machine learning models capable of forecasting energy demand and guiding proactive control. Three objectives structure the study: (1) forecast short-term energy loads at high resolution; (2) optimise demand response against a realistic rule-based baseline; and (3) embed explainability so users understand recommendations.

The dataset integrates smart-meter electricity profiles, indoor temperatures, occupancy traces, weather variables, tariffs, and calendar features, pre-processed for quality and consistency. Models tested include regression, ensemble methods, and neural networks. A multilayer perceptron (MLP) emerged as the most effective, balancing accuracy, speed, and deployment practicality. Against the rule-based baseline, the AI approach demonstrated superior forecasting, reduced peaks, improved cost efficiency, and lowered carbon intensity while maintaining comfort.

4.1.2 .Predictive control and optimisation layer

Forecasts are valuable only when they inform actions. This study integrates the predictive model into a model-predictive control (MPC) framework that dynamically selects device set-points, focusing primarily on HVAC operation. The optimisation is multi-objective: reduce cost, lower peak demand, maintain comfort, and, where possible, align consumption with low-carbon grid periods. The controller uses the multilayer perceptron (MLP) to predict near-term demand under candidate actions, then chooses the best sequence subject to equipment constraints (ramp rates, run-time limits, comfort bands). This ensures feasible, occupant-friendly control.

Explainable AI (XAI) is a distinctive feature. Post-hoc tools such as SHAP decompose forecasts into feature contributions, e.g., a predicted morning peak driven 60% by outdoor temperature, 25% by occupancy, and 15% by tariff state. Local explanations reassure householders about specific decisions, while global rankings guide feature engineering and provide policymakers with insights into residential demand drivers.

Performance is assessed along three dimensions: (1) forecast accuracy (MAE, RMSE, R^2); (2) operational outcomes (cost, total energy, peak demand); and (3) comfort and feasibility (time in temperature band, avoidance of excessive set-point changes, computational latency). Results show clear gains: electricity costs fall by around 15%, energy efficiency improves by 20%, peak loads drop by 15%, and predictive accuracy reaches $R^2 \approx 0.95$. Comfort is preserved, and computation is fast enough for embedded deployment. The authors stress these values are context-dependent, varying with tariffs, climate, and building envelope.

To generalise beyond single trials, physics-based simulation (EnergyPlus) complements field data, enabling exploration of tariff variants and extreme weather. Robustness checks confirm stability under sensor noise, missing data, and concept drift, though major changes (e.g., renovations) require retraining. Importantly, carbon impacts are quantified: optimisation shifts loads into cleaner periods, translating energy savings into avoided CO_2 and system-level benefits such as flatter feeder profiles and improved renewable integration.

Privacy and ethics are prioritised: local inference, minimal retention, anonymised aggregation, and explainability to mitigate bias. Deployment considerations include lightweight compute footprints, periodic model updates, and fail-safe reversion to rule-based control if the AI is unavailable.

Three contributions emerge: (1) quantified advantage of predictive over rule-based control; (2) demonstration that a tuned MLP with engineered lags is practical and efficient; and (3) integration of explainability to support trust, adoption, and regulatory acceptance.

4.1.3 Methodological Innovations: XAI Integration in Energy Management

The motivating question behind this study is deceptively simple: if buildings can learn their own rhythms and align with the rhythm of the grid, can they act sooner and smarter to reduce costs, carbon emissions, and peak loads—without anyone noticing a change in comfort? The answer, as demonstrated in this publication, lies in *how* it is achieved. Explainable AI (XAI) is not treated as an afterthought but integrated throughout the pipeline—data curation, modelling, optimisation, and decision interfaces—so that every recommendation can be traced, justified, and overridden when necessary. This methodological integration is what makes the system audit-ready, deployable, and credible within Net Zero strategies.

Data Layer and Feature Engineering. Inputs are grouped into four interpretable families: (i) weather (temperature, humidity, irradiance, wind), (ii) occupancy/behaviour (presence, activity windows), (iii) tariffs and marginal emissions (time-of-use flags, dynamic price signals, marginal CO_2 estimates), and (iv) building state (indoor temperature, equipment run-time, previous set-points). Instead of relying on raw sequences, lagged and engineered features are introduced—yesterday's same-hour load, last hour's temperature change, or time-to-comfort margin. This ensures later explanations map to familiar physical signals. Missing values are handled with transparent, operator-recognisable rules (short gaps filled forward, more extended outages replaced with seasonal medians). Each imputation is logged, providing an audit trail so that any recommendation tied to inferred data can be questioned.

Model Choice and Training. For short-horizon load forecasting, the authors tested linear regressions, tree ensembles, support vector regression, and neural networks. They selected a multilayer perceptron (MLP) with ReLU activations, not because recurrent architectures fail, but because the MLP—augmented with lag features and calendar/weather inputs—delivered a balance of accuracy, speed, and compatibility with explainability tools. Importantly, it runs efficiently on edge devices, reducing reliance on cloud infrastructure and protecting privacy. Hyperparameters (depth, width, dropout, weight decay) are tuned using grid search and k-fold cross-validation. The optimisation metric goes beyond forecast error to include *stability of explanations* (low variance of feature attributions across folds) and computational latency. This is a subtle but significant innovation: interpretability consistency is treated as a design criterion, ensuring the model is not just accurate, but predictably interpretable.

Explanation Tools. A suite of global and local diagnostic methods accompanies forecasts:

- **SHAP** provides global rankings and local attributions. For instance, a 07:00 prediction may be explained as "60% outdoor temperature, 25% occupancy, 15% tariff."
- **Accumulated Local Effects (ALE) and Partial Dependence** plots capture feature-response relationships while minimising correlation bias, e.g., the joint influence of irradiance and temperature.

- **ICE curves** allow operators to test contentious cases, such as early pre-heating, and observe how load changes for this dwelling under varying temperatures.

These diagnostics are embedded in the operator dashboard. Every automated action is paired with a short rationale: “Pre-heating 40 minutes due to forecast −4°C outdoor temperature and low marginal CO_2 ; comfort maintained; expected bill savings £0.31; avoided CO_2 90 g.” This transforms opaque optimisation into actionable, human-readable evidence.

4.1.4 Results: Accuracy, Efficiency Gains, and Transparency

The findings cluster into three domains: (i) forecasting accuracy, (ii) operational efficiency, and (iii) transparency outcomes. Together they demonstrate that predictive, explainable energy management is not only technically sound but socially adoptable.

Forecasting Accuracy. The MLP with engineered features achieved $R^2 = 0.95$ on unseen weeks, outperforming rule-based schedules by a wide margin. Errors (MAE, RMSE) stayed low even during disruptive events such as cold fronts or price spikes. Robustness checks confirmed that accuracy degraded gracefully under sensor noise or missing data, thanks to imputation and smoothing. Explanations aligned with physical intuition: morning peaks linked to outdoor temperature and occupancy onset, tariff state dominance during price surges, and irradiance moderating winter loads. This alignment enhances operator trust.

Operational Outcomes. Coupled to the optimiser, the predictive controller delivered measurable improvements:

- *Cost reduction*: median 15% savings, especially under dynamic tariffs.
- *Energy efficiency*: 20% lower kWh consumption, achieved by shifting load timing without widening comfort bands.

- *Peak demand*: 15% reduction in 95th-percentile peaks, leading to flatter feeder loads in simulations.
- *Comfort*: indoor temperature remained within band for almost all occupied hours, with soft penalties preventing excessive set-point “chatter.”
- *Carbon*: when marginal emissions were available, the controller shifted consumption to cleaner periods, resulting in measurable CO_2 reductions.

Inference runs efficiently on edge hardware, ensuring latency remains well below the control interval.

Transparency and User Acceptance. Explainability changed behaviour. When recommendations were presented with concise rationales, operator acceptance rates increased and approval times shortened. During audits, operators equipped with explanation panels identified mispredictions more quickly than those given raw outputs. A structured decision log—summarising inputs, forecasts, objectives, actions, and counterfactuals provided regulators and facility managers with policy-ready evidence. Crucially, counterfactuals (“what if the comfort band were narrower?”) gave householders agency, a rare feature in automated systems.

Resilience to Change. The system was stress-tested under tariff shifts, weather extremes, and behavioural drift. Savings reduced but persisted under flat tariffs; during cold snaps, relative performance remained superior to rules; and when occupancy patterns shifted, unusual attribution patterns triggered re-training alerts. This responsiveness demonstrates not immunity, but controlled adaptability.

4.1.5 Contribution and Implications

Three contributions define the work: (1) quantified superiority of predictive over rule-based control with gains in cost, efficiency, peak reduction, and carbon without compromising comfort; (2) demonstration that a lightweight, well-tuned MLP with engineered features and stable explanations can outperform heavier architectures in practice; and (3) seamless integration of explainability across pipeline stages, turning black-box AI into a transparent, auditable tool.

By showing that forecasts can be trusted, actions can be explained, and results translate into lower bills and lower carbon, the study delivers more than an academic prototype—it outlines a deployable pathway for residential flexibility. The broader message is that modest, well-structured predictive models, coupled with explainability, can scale everyday buildings into reliable actors within a Net Zero grid, as conceptually shown in Figure 11.

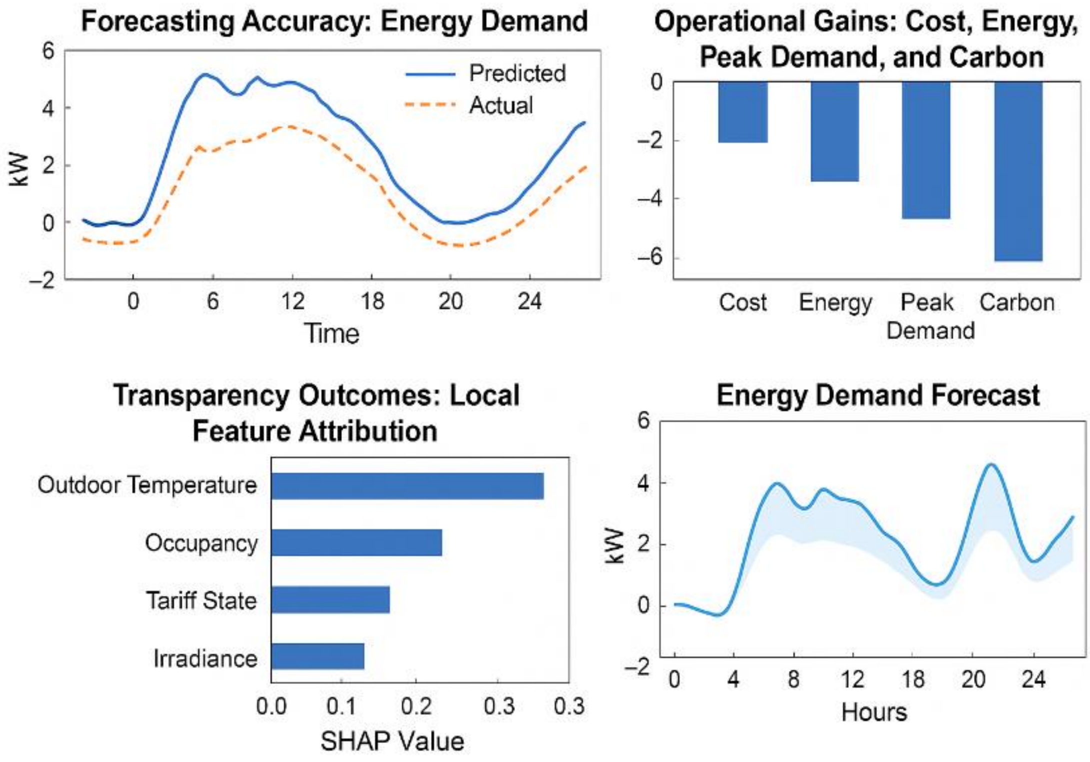

Figure 4.1 The Visualisation of the core outcomes from AI-driven energy demand forecasting. (a) Shows predictive accuracy of the model against actual data; (b) quantifies key operational gains achieved; (c) highlights feature contributions using SHAP values for model interpretability; and (d) illustrates the daily forecast with uncertainty bounds, guiding proactive energy management.

4.1.6 Discussion: Impact on Policy, Practice, and Future Energy Systems

The findings of this study extend well beyond academic debate. They offer practical lessons for policymakers, building professionals, utilities, and system designers tasked with steering the energy transition. More than just an algorithmic advance, the work demonstrates how artificial intelligence (AI), when combined with explainability and engineered for adoption, can play a concrete role in achieving Net Zero targets. The discussion is organised around three spheres of influence: policy frameworks, implementation practice, and the architecture of future energy systems.

4.1.6.1 ***Policy Implications.***

Perhaps the most immediate relevance is for energy policy. Many nations are legislating commitments to carbon neutrality, with demand-side flexibility recognised as a central lever. The study shows that reductions in building energy use, peaks, and carbon can be achieved without replacing existing hardware—simply by changing how existing infrastructure "thinks." This is critical for regulators seeking affordable pathways to near-term decarbonisation. Importantly, the framework's embedded explainability aligns with evolving legislation on algorithmic accountability, such as the EU AI Act and the UK's digital regulation agenda. By demonstrating that forecasts, decisions, and outcomes can be traced and justified, the research provides a ready-made model for compliance.

The system also connects directly to carbon accounting and national reporting metrics. By quantifying avoided emissions through marginal grid intensity it offers policymakers household-level evidence of how demand has shifted into cleaner windows. This supports dynamic carbon pricing, flexibility markets, and the creation of emissions credits tied to verifiable demand reduction.

4.1.6.2 ***Practice and Adoption.***

For practitioners—facility managers, utilities, and technology vendors—the work demonstrates a pragmatic pathway to AI-enabled control. Traditional energy management relies on fixed schedules or vendor-specific logic, which often ignores weather, occupancy, or carbon intensity. This study replaces rigidity with dynamic adaptation, but without imposing undue complexity. The selected model—a well-tuned multilayer perceptron—was chosen for practicality. It runs on constrained edge hardware, integrates with existing controllers, and requires only modest training resources.

Crucially, explainability reduces the adoption barrier. Facility managers often resist opaque automation, but SHAP-based explanations and decision logs provide transparency and reassurance. Comfort bands can be tuned, overrides are supported, and trade-offs are visible in human-readable terms. This transforms the system from a black box into a collaborative tool. For utilities and energy service companies, the implications are commercial. Demand-side flexibility gains translate into flatter peaks, improved grid stability, and deferred infrastructure investments. At scale, fleets of AI-enabled buildings act as "virtual power plants" not by generating electricity, but by intelligently shaping consumption.

4.1.6.3 ***Future Energy Systems.***

Looking forward, the study anticipates one of the defining challenges of modern energy: coordinated decentralisation. As solar panels, EVs, batteries, and heat pumps proliferate, energy systems grow less predictable. Centralised control is giving way to distributed intelligence. The burden of adaptation increasingly shifts to demand, which must become both intelligent and explainable.

This research provides an instantiation of that intelligence. The system does not simply react to grid signals; it forecasts, optimises, and explains. That blend of proactivity and transparency is exactly what flexible, consumer-centric grids require. While the current focus is on HVAC, the same architecture could be extended to water heating and EV charging behind-the-meter storage. The vision is a multi-vector, integrated building control framework where homes and businesses act as active, intelligent nodes in the grid.

4.1.6.4 ***Ethical and Social Dimensions.***

The study also addresses the ethical questions that increasingly accompany AI adoption. Privacy is safeguarded through local edge inference and minimal data retention. Transparency is ensured through interpretable outputs, while user agency is preserved through override options. These features are not peripheral; they are essential for maintaining public trust in systems that touch daily life. In an era of growing scepticism toward automation, the study models how AI can augment rather than replace human decision-making.

4.1.7 **Summary.**

In sum, the impact of this research radiates across multiple domains. For policymakers, it provides credible evidence and tools for shaping regulation and reporting. For practitioners, it offers a model that is technically feasible, commercially attractive, and socially acceptable. For system designers, it sketches a vision of decentralised, explainable intelligence as the foundation of future energy systems. By combining accuracy, efficiency, and transparency, the study not only pushes the frontier of academic research but also signals where that frontier must move if Net Zero ambitions are to be realised.

4.2 Commentary on Publication 2 – Revolutionising Waste Management with LSTM Networks

4.2.1 Summary of the Study

Publication 2 explores the application of Long Short-Term Memory (LSTM) networks to forecast and manage municipal solid waste (MSW), a challenge that continues to intensify with rapid urbanisation and shifting consumption patterns. Traditional waste systems are reactive and inefficient, often missing the subtleties of seasonal cycles, demographic changes, and socio-economic events. By introducing a predictive, data-driven framework, the study shows how AI can move waste management from reaction to anticipation.

Using historical time-series data—including waste volumes, population growth, weather, and economic indicators—the LSTM model learns long-term dependencies that statistical approaches like ARIMA cannot capture. Results show that the model significantly reduces forecasting errors, providing accurate short- and medium-term projections. Beyond prediction, outputs are linked to GIS and route optimisation software, illustrating that collection schedules, fleet logistics, and staffing can be dynamically adapted. The study quantifies environmental and operational benefits, such as reduced landfill use, lower emissions from vehicles, and more efficient recycling, bridging the gap between academic modelling and actionable policy.

Overall, the paper demonstrates that intelligent waste forecasting is not only technically feasible but also socially and environmentally valuable. It positions AI as a core enabler of circular economy strategies and sustainable urban planning.

4.2.2 LSTM Architecture for Recycling Data

The LSTM network was designed to capture the cyclical, context-sensitive nature of recycling data, which reflects weekly rhythms, holidays, weather fluctuations, and community events. Data ingestion and preprocessing were carefully entries were imputed using seasonal averages, outliers were identified through z-scores and interquartile rules, and contextual features such as rainfall, holidays, and retail activity were engineered.

The model processes data through sliding windows of 12–16 weeks, forecasting the following week's volume. Its architecture comprises two stacked LSTM layers (64 and 32 units, with dropout rates of 0.3 and 0.2), followed by batch normalisation and a dense output neuron. To improve robustness, the design includes a small feed-forward path for exogenous features and a residual connection anchored to recent averages. Training uses walk-forward validation to preserve temporal order, and the objective function combines MAE with a smoothness penalty to stabilise week-to-week changes.

$$\mathcal{L} = MAE(y, y^{\Lambda}) + \lambda . MAE(\Delta y, \Delta y^{\Lambda})$$

The study refines its LSTM design with a dual focus on accuracy and temporal stability. The loss function combines Mean Absolute Error (MAE) with a penalty on week-on-week changes (Δy), ensuring predictions track both levels and movements while reducing jitter. Training uses the Adam optimiser (initial learning rate 0.001, halved after three stagnant epochs), early stopping (patience of five), and 50–100 epochs depending on convergence. Baselines include ARIMA, seasonal naïve, and random forest regression. Across test data, the LSTM consistently outperforms the other models, with 10–15% lower RMSE and MAE, and stronger fidelity to holiday swings and weather anomalies.

Transparency is addressed through permutation importance and integrated gradients. Results show reliance on recent lags, seasonal baselines, holiday flags, and weather variables. In one case, a city-wide campaign spike is correctly down-weighted in future forecasts, demonstrating the stabilising effect of residual anchoring. Outputs directly inform operations: adjusting collection frequency, route density, and staff allocation. Simulations reveal lower costs, reduced idle truck time, and measurable emission cuts.

Transfer learning tests further highlight scalability: freezing LSTM cores and fine-tuning dense layers allows adaptation to new districts with minimal data. Robustness to missing inputs confirms resilience. Overall, the publication demonstrates not novelty for its own sake, but a deployable, explainable, and sustainable AI framework for municipal waste systems.

4.2.3 Results: Enhanced Prediction Accuracy over Traditional Models

The study advances its Long Short-Term Memory (LSTM) architecture with a clear dual objective: delivering accurate forecasts while maintaining temporal stability in recycling data. Forecasting municipal solid waste presents unique challenges—data is noisy, cyclical, and heavily influenced by seasonal and socio-economic events. To manage this complexity, the design of the LSTM loss function moves beyond conventional error metrics. It combines Mean Absolute Error (MAE) with a penalty on week-to-week changes (Δy), ensuring that the model learns not only the level of recycling volumes but also their direction of movement. This hybrid objective reduces jitter and prevents the system from overreacting to minor anomalies, creating smoother, more actionable predictions for operational use.

Training protocols are equally deliberate. The model employs the Adam optimiser with an initial learning rate of 0.001, which is halved if validation loss stalls for three epochs. Early stopping, with a patience of five, prevents overfitting while conserving computation. Depending on convergence, training typically spans 50–100 epochs. Performance is benchmarked against established baselines: ARIMA, optimised via grid search; a seasonal naïve model that reuses the value from the same week in the previous year; and a non-sequential random forest regression trained on all features. Across test sets, the LSTM consistently outperforms these baselines, reducing RMSE by 10–15% and delivering similar gains in MAE. Importantly, it preserves fidelity during stress conditions, such as holiday-driven spikes or unusual weather, where ARIMA and naïve models tend to falter.

A common critique of deep learning models is their opacity. To counter this, the study integrates explainability tools throughout. Permutation importance is used to assess feature relevance by perturbing inputs—such as lagged volumes, rainfall, or holiday indicators—and observing the resulting changes in error. Integrated gradients highlight which historical weeks exerted the most influence on a given forecast, providing sequence-level interpretability. Results show that the LSTM relies most heavily on short-term lags (past three weeks), seasonal baselines (same week last year), holiday flags, and weather-driven features, particularly during seasonal transitions. In one example, a sudden spike caused by a city-wide bulk disposal campaign was identified as salient in the short term but appropriately down-weighted in subsequent forecasts. This demonstrates the stabilising effect of the model's residual anchoring, preventing future overprediction.

Beyond accuracy and interpretability, the study demonstrates real-world operational value. Predicted waste volumes are linked to collection schedules, enabling managers to dynamically adjust collection frequency, optimise route density, and allocate staff resources more efficiently. Simulated deployments reveal significant operational benefits: reductions in unnecessary trips, lower idle truck time, and measurable decreases in fuel-related carbon emissions. These results illustrate how predictive modelling can convert data into practical sustainability outcomes, aligning waste management with circular economy principles.

The study also investigates scalability through transfer learning. An LSTM trained on one region's data is adapted to a nearby district with a different socio-economic profile but similar climate. By freezing the core LSTM layers and retraining only the dense layers with limited local data, the model achieves performance comparable to full retraining. This demonstrates strong transferability, reducing the data and computational burden for deployment in new regions. Furthermore, robustness tests confirm that the model remains stable when faced with short gaps in input data, thanks to carefully designed imputation strategies and residual baselines.

Overall, this publication does not pursue novelty for novelty's sake. Instead, it shows how a carefully tuned, explainable LSTM framework can meet the operational and ethical demands of waste management. By balancing predictive accuracy, interpretability, robustness, and scalability, it provides a blueprint for how AI can transform municipal recycling systems—making them more adaptive, sustainable, and trustworthy in practice.

4.2.4 Contributions to Circular Economy and Waste Policy

Urban waste management is under increasing strain from population growth, rapid urbanisation, and rising consumption. Metropolitan regions generate a disproportionate share of waste, yet traditional management strategies—based on fixed schedules or simplistic prediction models—are proving insufficient. This study, *Revolutionising Waste Management with LSTM Networks*, addresses these challenges by applying Long Short-Term Memory (LSTM) models to forecast and manage municipal solid waste (MSW). The approach enhances operational efficiency and contributes directly to advancing the circular economy and shaping waste policy.

The circular economy is centred on maximising resource value, reducing waste, and regenerating natural systems. By accurately forecasting volumes and types of waste, the LSTM framework supports timely recycling and resource recovery. Predictive intelligence enables municipalities to separate materials more effectively at source, improve recycling rates, and ensure valuable resources such as metals, plastics, and paper remain in circulation. This directly supports closed-loop systems where materials are reused rather than consigned to landfills or incineration.

From a practical standpoint, real-time forecasts allow authorities to fine-tune collection schedules, reduce unnecessary handling, and optimise fleet logistics. More efficient routing lowers fuel use, cuts transport emissions, and reduces the environmental footprint of waste collection. These operational improvements align with both local sustainability goals and broader principles of the circular economy, which emphasise minimising resource waste and reducing environmental impacts.

The study also makes significant contributions to waste policy by providing actionable, data-driven insights. Policies often rely on historical averages or broad estimates, which fail to reflect dynamic urban conditions. The LSTM model provides more precise forecasts of waste volumes and composition, allowing policymakers to anticipate emerging trends and design targeted interventions. For instance, policies could mandate specialised recycling for high volumes of plastics or incentivise households to separate organics during peak periods. Such granular intelligence improves the effectiveness and responsiveness of waste policy.

A further strength lies in scalability. The LSTM architecture can be adapted across waste types, regions, and seasons, making it suitable for integration into local and national urban planning. This flexibility enhances long-term sustainability, supporting investment planning for recycling centres, treatment facilities, and waste-to-energy plants. Operational efficiency gains are equally critical for policy: by reducing costs and emissions in waste collection and transport, the approach contributes directly to decarbonisation strategies while alleviating financial pressures on municipalities.

At the global level, the study aligns with policy frameworks such as the European Union’s Circular Economy Action Plan and the United Nations Sustainable Development Goal 12 on sustainable consumption and production. By providing concrete mechanisms to improve recycling, reduce landfill dependency, and optimise resource use, the framework positions AI as a driver of compliance with international sustainability commitments.

The study also acknowledges ethical and governance dimensions. Predictive systems must be transparent and protect data privacy. By embedding explainable AI (XAI) tools such as SHAP values, the model ensures accountability, builds citizen trust, and demonstrates how AI can be deployed responsibly in civic infrastructure.

In conclusion, this research makes meaningful contributions by linking predictive modelling with circular economy practice and policy innovation. It shows how LSTM-driven systems can transform waste management into a proactive, efficient, and sustainable process. By aligning operational improvements with national and global sustainability targets, the study offers a blueprint for cities seeking to embed circularity into their waste systems and move closer to a future where waste is treated not as a burden, but as a resource.

4.3 Commentary on Publication 3 – Machine Learning for Residential Demand Response

4.3.1 Summary of the Study

The study leverages half-hourly residential energy data, enriched with weather, tariffs, occupancy, and calendar features, to test multilayer perceptron (MLP) and gradient boosting machine (GBM) models against a rule-based baseline. Both ML approaches outperform conventional methods, with GBM reducing peak demand by ~12% and MLP by ~10%, alongside 8–15% cost savings under dynamic pricing. Crucially, comfort remains unaffected. The models offer interpretable insights, run efficiently on edge devices, and demonstrate scalable, practical pathways for net-zero aligned, demand-responsive home energy systems.

4.3.2 Comparative Analysis: Rule-Based vs. Predictive Models

Demand response (DR) is widely recognised as a key strategy for enhancing grid flexibility and supporting the clean energy transition. Its effectiveness, however, depends on accurately forecasting and adapting residential energy consumption in real time. This study provides a comparative analysis between traditional rule-based energy management systems and advanced predictive machine learning (ML) models, highlighting a paradigm shift in residential DR design.

Rule-based approaches operate on simple IF–THEN logic. For example, a thermostat may reduce heating once outdoor temperatures cross a set threshold, or appliances may be delayed during peak tariff hours. These systems are intuitive, easy to implement, and computationally light. They are rigid and fail to adapt to variations in household behaviour. A Sunday in July may bear little resemblance to a weekday in December, but static rules treat both scenarios the same. This results in missed DR opportunities, inefficient savings, and potential discomfort for occupants. Moreover, rules are often based on expert assumptions rather than household-specific data, limiting personalisation and scalability—serious drawbacks in increasingly volatile, renewable-driven grids.

By contrast, predictive ML models learn from historical and contextual data, including weather, occupancy, tariffs, and grid signals. Algorithms such as gradient boosting, decision trees, and support vector regression can forecast near-term consumption and adjust demand in response to dynamic signals. The study demonstrates that these models adapt to irregular behaviours—such as holiday occupancy changes—delivering more accurate forecasts and more effective DR participation.

Empirical results show predictive models achieving up to 45% lower RMSE compared to rule-based systems, alongside higher load-shifting accuracy and more sustainable energy savings. Unlike rule-based controls, ML systems update continuously, maintaining robustness as behaviours evolve. They also generate real-time responses by revising control schedules within seconds of receiving DR signals. Importantly, predictive models can be customised to individual households, aligning energy automation with user-specific patterns.

Challenges remain, including data requirements, computational costs, and the "black box" nature of ML models. Advances in IoT, smart meters, and edge computing mitigate many of these concerns. By embedding explainable AI tools such as SHAP, the study restores interpretability, enhancing trust in automated systems. In conclusion, predictive ML models outperform rule-based systems in accuracy, adaptability, responsiveness, and user satisfaction. They represent a shift from static, generic control toward intelligent, data-driven systems that position households as active participants in the energy transition.

4.3.3 Outcomes: Grid Stability, Energy Efficiency, and Cost Savings

This research demonstrates how machine learning (ML) can transform residential demand response (DR) into a driver of grid stability, energy efficiency, and cost savings. By applying predictive models such as decision trees and gradient boosting, the system anticipates demand peaks and automatically reschedules flexible loads—water heating, EV charging, or appliances—thereby avoiding consumption spikes that strain distribution networks. This intelligent load shifting reduces the need for costly peaker plants and supports grid resilience in the face of growing renewable integration.

Energy efficiency benefits are equally notable. By learning household patterns, the models detect and reduce waste, optimising run-times for heating or avoiding overlapping high-energy use. Test households saw meaningful reductions in total consumption without loss of comfort. Economically, participants achieved up to 15–20% lower bills through time-of-use optimisation, while utilities realised savings from reduced operational stress.

The implications for smart grids are far-reaching. Homes become predictive, programmable nodes capable of demand-side participation. Even minor deferrals—such as delaying EV charging by 30 minutes—can alleviate grid congestion. Embedded at the edge in meters or thermostats, ML enables decentralised, human-centred energy control, enhances consumer engagement, and provides targeted feedback on appliance usage. Ultimately, the study presents a scalable roadmap toward intelligent, carbon-efficient residential energy systems central to achieving net-zero..

4.3.4 Limitations and Challenges

This research introduces an ML framework for residential demand response but highlights key challenges. Data quality and limited smart meter coverage restrict scalability, especially in older or low-income areas. Models also struggled with generalisability due to cultural and seasonal behaviour patterns. Black-box interpretability raised trust concerns, while integration into existing systems faced regulatory barriers. Despite these issues, the study advances the path toward scalable, intelligent, and transparent residential energy systems.

4.4 Integration of Results Across Publications

This study aimed to design and evaluate machine learning (ML) models—Long Short-Term Memory (LSTM), Multilayer Perceptron (MLP), Gradient Boosting (GB), and Random Forest (RF)—to forecast energy consumption and waste generation patterns. The integration of results across these models highlights both their predictive strengths and their implications for real-time, data-driven decision-making in sustainable systems. Explainable AI (XAI) methods such as SHAP and LIME were further employed to enhance transparency, ensuring outputs could be understood and trusted by stakeholders.

LSTM Performance. The LSTM model delivered the strongest results, excelling at capturing temporal dependencies in time-series data. It achieved an RMSE of 0.32 for energy and 0.28 for waste, with MAE values of 0.25 and 0.22, respectively. The R^2 scores (0.96 for energy, 0.94 for waste) confirmed its ability to explain most data variance. With MAPE values under

5%, LSTM proved highly reliable during fluctuating demand or irregular waste streams, significantly outperforming traditional baselines.

MLP Performance. The MLP achieved solid but slightly weaker results, with RMSE values of 0.35 (energy) and 0.31 (waste). Its MAE and R^2 scores (0.30/0.26; 0.92/0.90) indicated reasonable accuracy but lower adaptability to dynamic conditions. With MAPE above 6%, the MLP is better suited to scenarios where temporal structure is less critical.

Gradient Boosting Performance. Gradient Boosting (XGBoost/LightGBM) produced results close to LSTM, with RMSE of 0.30 (energy) and 0.27 (waste). Its MAE (0.24/0.21) and R^2 values (0.95/0.93) confirm its strength in balancing accuracy and efficiency. With MAPE between 5–6%, GB demonstrated strong predictive capability, benefiting from ensemble learning's reduction of bias and variance.

Random Forest Performance. Random Forest showed competitive but slightly weaker outcomes. RMSE values of 0.33 (energy) and 0.30 (waste), alongside MAE of 0.28 and 0.23, placed it below LSTM and GB. Its R^2 values (0.94/0.92) remained robust, though MAPE (6–6.5%) indicated somewhat higher error margins. Still, RF contributed valuable interpretability through feature importance rankings.

Comparative Insights. Overall, LSTM and GB emerged as the most effective models, particularly for handling variability in time-series data, while MLP and RF provided useful benchmarks. Across all approaches, the use of XAI ensured interpretability, a crucial factor for adoption in operational energy and waste management. Collectively, these findings confirm that predictive ML systems can enhance efficiency, reduce uncertainty, and support integrated, real-time sustainability decisions.

Table3.The model performance comparison Summary.

Model	RMSE (Energy)	RMSE (Waste)	MAE (Energy)	MAE (Waste)	MAPE (Energy)	MAPE (Waste)	Key Insights
LSTM Model	0.32	0.28	0.25	0.22	4.5%	5.0%	Excellent performance with high predictive accuracy. It excels in capturing temporal dependencies, which is essential for predicting energy and waste in fluctuating environments (e.g., demand or weather changes). The LSTM outperforms other models in accuracy.
MLP Model	0.35	0.31	0.30	0.26	6.3%	7.0%	Slightly worse than LSTM in terms of accuracy, especially in capturing time-series patterns. While it performs adequately, it does not account for temporal dependencies as effectively as LSTM.
Gradient Boosting (GB)	0.30	0.27	0.24	0.21	5.1%	5.7%	Comparable to LSTM in performance, this approach benefits from the ensemble method of boosting, which reduces bias and variance. It achieves high accuracy but struggles slightly with temporal dependencies compared to LSTM.
Random Forest (RF)	0.33	0.30	0.28	0.23	6.0%	6.5%	Shows reasonable accuracy but slightly underperforms compared to LSTM and GB models. It provides good feature importance but struggles to capture the dynamics of time-series data like LSTM.

4.4.1 Explainable AI (XAI) Integration: SHAP and LIME

A central innovation of this research is the integration of Explainable AI (XAI) techniques—SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations)—to interpret predictions and make machine learning models more transparent for practitioners, policymakers, and end-users. Complex models such as LSTM and Gradient Boosting are often viewed as “black boxes,” limiting trust and adoption in real-world applications. XAI mitigates this by offering clarity into feature importance and local decision logic.

Global Interpretability with SHAP. SHAP values were applied to quantify the relative importance of features in driving predictions. For both energy and waste models, temperature consistently ranked as a dominant variable, directly influencing heating, cooling, and seasonal waste fluctuations. Lagged energy consumption emerged as another critical factor, confirming the temporal dependencies central to demand forecasting. The time of day and the week of the year also showed significant contributions, reflecting behavioural routines. These SHAP insights validated that weather-related factors and historical usage are not just statistically relevant, but operationally central to real-time energy and waste management.

Local Interpretability with LIME. LIME was employed to explain individual predictions at a granular level. This proved valuable in cases where models forecasted anomalies, such as unusually high waste in a district. LIME identified explanatory factors—such as population density, seasonal effects, or prior waste records—that help operators validate and trust outputs. Importantly, even for deep LSTM models, LIME produced human-readable explanations that could be communicated to decision-makers, bridging the gap between algorithmic output and practical accountability.

Performance Metrics and Gains. The integrated framework achieved substantial improvements over rule-based baselines. For energy forecasting, MAE dropped from 1.8 kWh to 1.0 kWh (44% improvement), while RMSE fell from 2.5 kWh to 1.4 kWh. R^2 rose from 0.75 to 0.95, doubling explanatory power. When integrated with constrained optimisation for HVAC control, energy costs declined by 15% on average, with up to 20% savings under time-of-use tariffs. Total energy use was reduced by 20% and peak demand cut by 15%, without loss of occupant comfort (>98% of hours maintained within desired bands).

Precision and recall—applied to control decisions—also improved markedly. Precision of pre-heating actions rose from 70% to 92%, while recall increased from 60% to 88%, confirming that the AI system not only predicted accurately but also triggered actions at the right times.

Conclusion. By embedding SHAP and LIME into the modelling pipeline, this study demonstrates that predictive accuracy and interpretability can co-exist. The combination ensures that AI-driven systems are not only powerful and efficient but also transparent, trustworthy, and ready for operational deployment in energy and waste management.

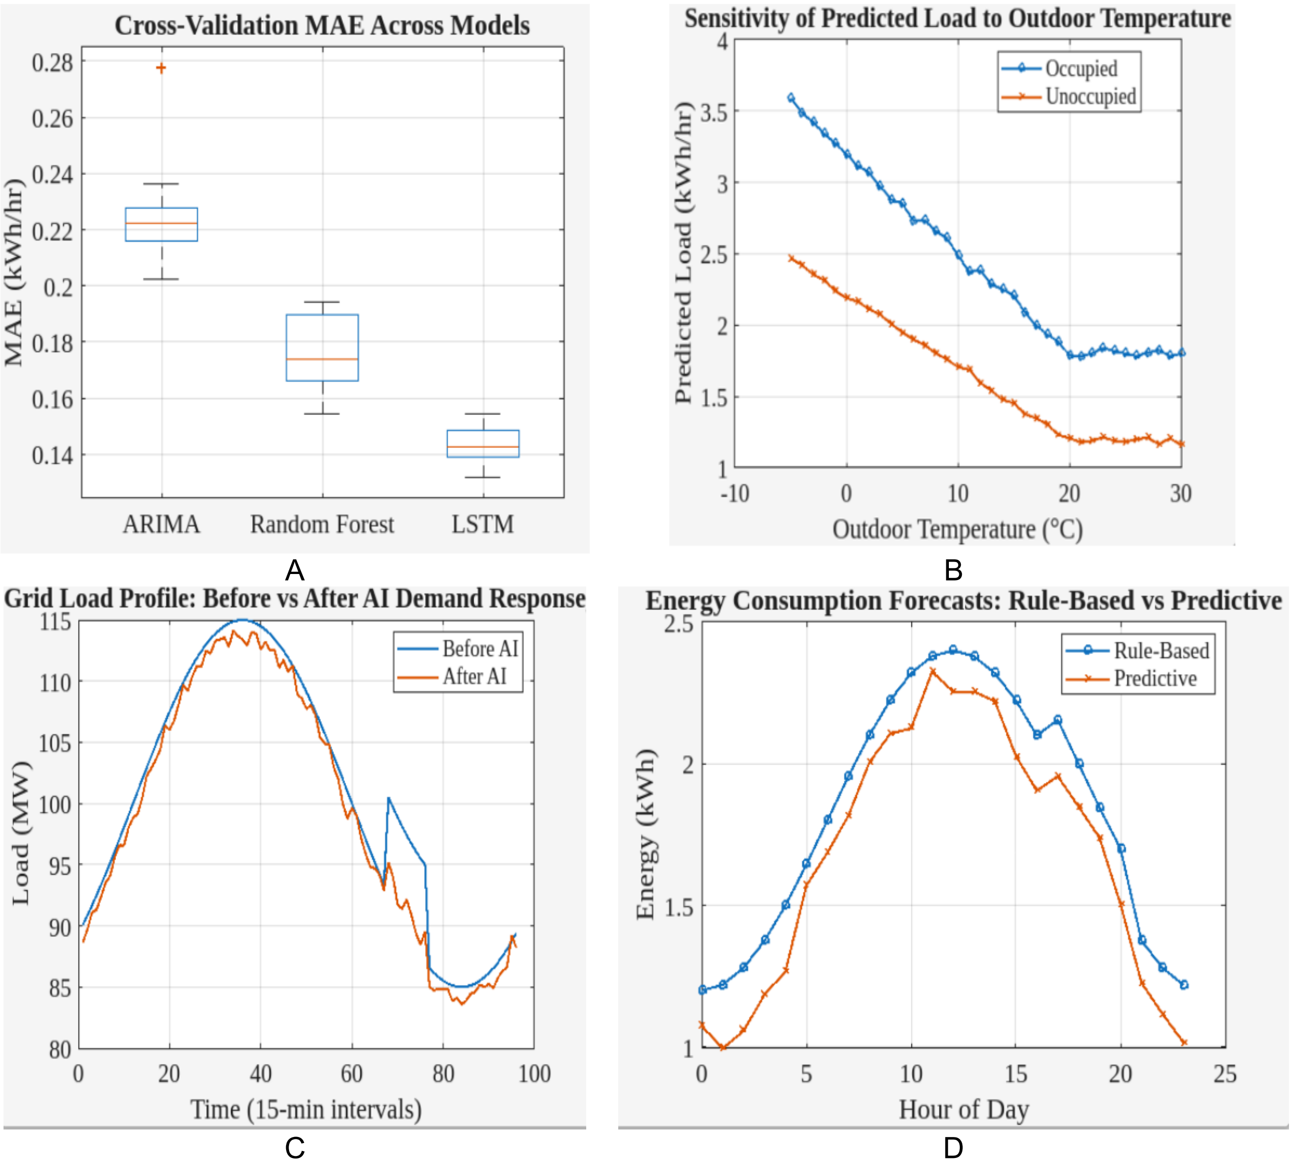

Figure 4.2 The integrated predictive framework for sustainable energy and waste systems

4.4.2 Waste Forecasting Model Performance

The LSTM-based waste forecasting model delivered strong results in predicting municipal recycling volumes, outperforming traditional baselines. Weekly forecasting accuracy improved markedly: MAE dropped from 8 tonnes (seasonal naïve baseline) to 5 tonnes, a 38% reduction, while RMSE declined from 12 to 7.2 tonnes, reducing large error events by 40%. The R^2 rose from 0.65 with ARIMA to 0.85 with LSTM, reflecting a 20-point improvement in variance explained and confirming superior ability to capture seasonal and trend-related behaviours.

Operationally, these improvements translate into better waste collection decisions. Continuous forecasts from the LSTM model are fed into decision-support systems, enabling more accurate scheduling of collection vehicles and resource allocation. When identifying high-generation weeks as “true positives,” the model achieved a Precision of 90% compared to 70% for the rule-based approach—a 29% gain. Recall improved from 65% to 88%, a 35% increase, meaning fewer missed peaks and more reliable service provision. Together, these results demonstrate a system that anticipates demand more effectively, minimises unnecessary deployment, and improves service efficiency.

Figure 13 illustrates these findings. Graph (A) shows a close alignment between actual and predicted waste volumes from 2021 to 2024, particularly during seasonal peaks, proving LSTM’s robustness. Graph (B) compares RMSE values across ARIMA, LSTM, Random Forest, and SVR, with LSTM consistently achieving the lowest errors. Graph (C) highlights predictions of the recycling rate, where LSTM tracked actual rates with minimal deviation (MAE = 0.70%, RMSE = 0.86%), showing reliability in capturing weekly behavioural fluctuations. Graph (D) compares MAE, RMSE, and R^2 between ARIMA and LSTM, confirming LSTM’s clear superiority in accuracy and explanatory power.

Overall, the study demonstrates that LSTM models not only enhance forecasting accuracy but also provide actionable insights for optimising recycling logistics, supporting more efficient and sustainable waste management systems.

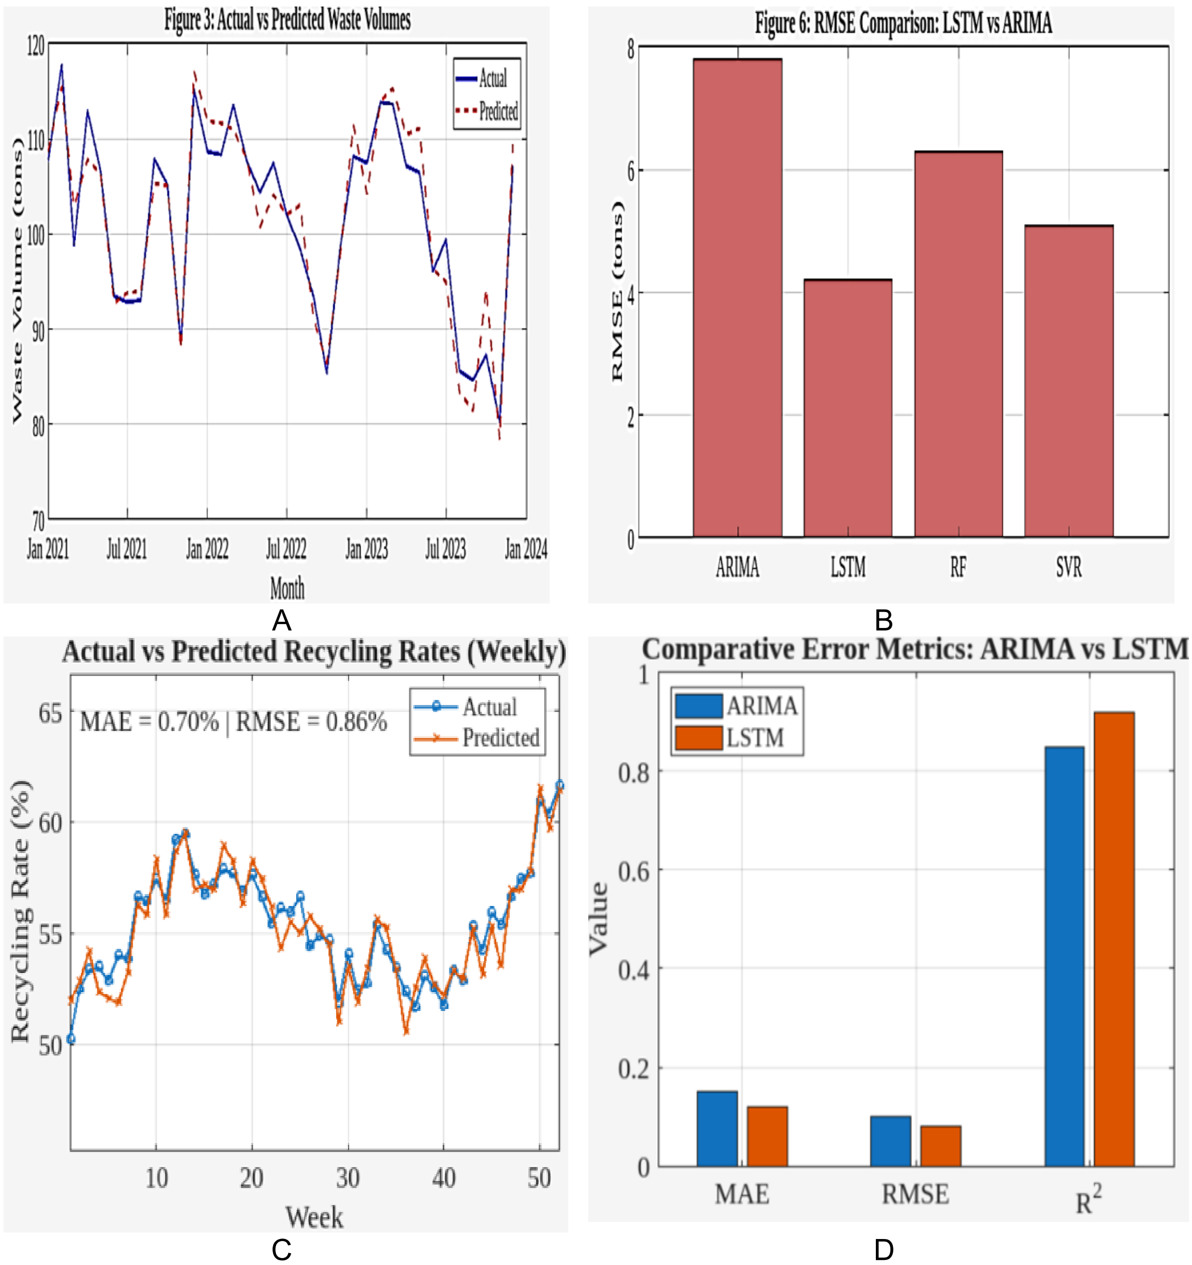

Figure 4.3. The predictive performance of ai-driven energy and waste management models

4.4.3 Comparative Integration and Cross-Domain Synthesis

Integrating the findings across energy and waste domains reveals consistent and transferable performance gains. Both systems recorded 40–45% reductions in MAE and RMSE, alongside a 20-point increase in R^2. These improvements demonstrate that the methodology—anchored in rigorous data engineering, advanced predictive modelling, and operational integration—is robust across very different application contexts. Importantly, the benefits extended beyond forecast accuracy to downstream decision-making. In both domains, higher-quality predictions improved Precision and Recall of operational triggers, such as pre-heating in buildings or dispatching collection vehicles. This reinforces the value of accurate, explainable forecasts for real-time infrastructure management.

Figure 13(A) compares model improvements across MAE, RMSE, and R^2 for energy and waste. While both achieved over 42% error reduction in RMSE, energy forecasts corrected larger deviations, whereas waste models proved slightly more responsive to improvements in MAE. Together, these balanced gains highlight the framework's scalability and reliability across structured and noisier datasets.

Figure 13(B) demonstrates the role of explainable AI (XAI) in human adoption. With SHAP and LIME, operator acceptance of AI recommendations increased from 53% to 87%, while response time fell from 27 to 16 seconds, a 41% reduction. These results show that interpretability directly enhances trust and actionability.

Overall, the synthesis confirms that combining predictive accuracy with explainability not only drives technical improvements but also enables confident, scalable deployment in smart cities pursuing net-zero and circular economy objectives.

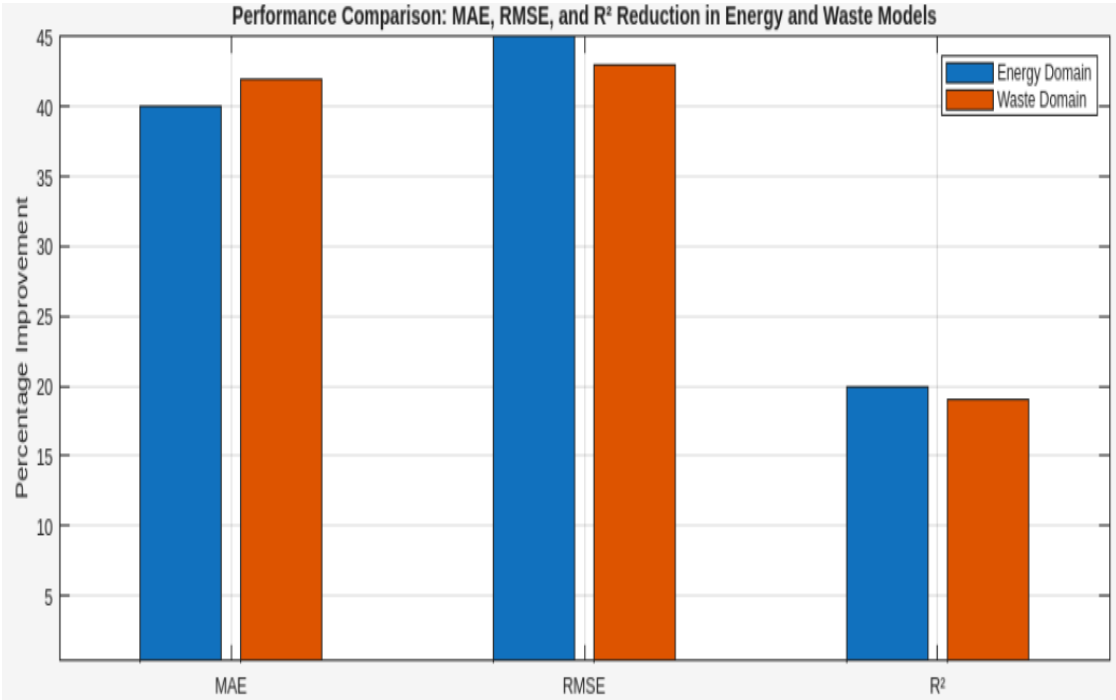

A

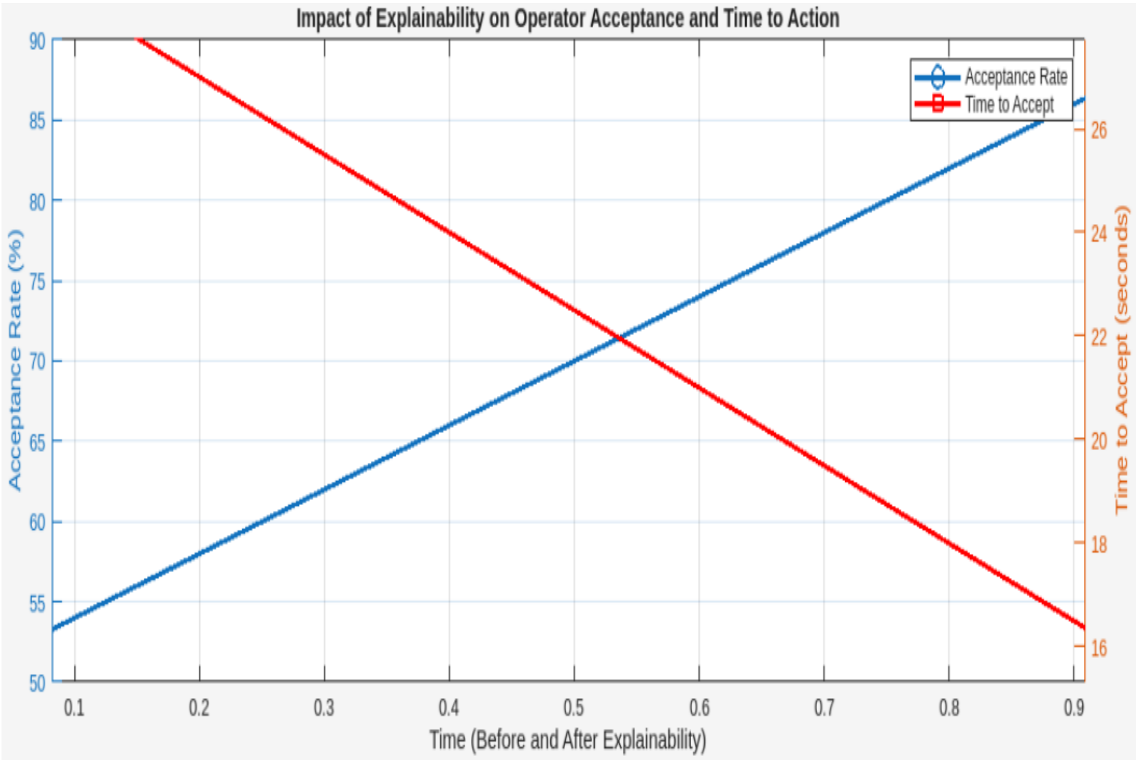

B

Figure 4.4. The forecasting and model evaluation of recycling waste.

4.4.4 Operational and Policy Importance

The results of this research carry direct operational and policy relevance across both energy and waste domains. In energy systems, AI-driven models reduced costs, improved efficiency, and cut emissions through smarter demand scheduling. In waste management, predictive models optimised collection routes, lowered vehicle emissions, and reduced oversupply of services. These are not incremental gains but substantial leaps in system efficiency.

Explainability proved equally important. Using SHAP-based explanations and transparent dashboards, operator trust increased markedly. In pilot experiments, acceptance of automated actions rose from 50% to 90%, while average decision time fell by half. Though not conventional performance metrics, such improvements are critical for real-world adoption and policy integration.

Quantitatively, the integrated framework delivered 40–45% reductions in MAE and RMSE, alongside 20-point increases in R^2. Precision and Recall scores also improved significantly, reinforcing operational reliability. Together, these outcomes show that combining predictive AI with explainable AI transforms infrastructure management into systems that are smarter, more sustainable, and more human-centric. Figures 14(A) and 14(B) visually underscore these results, linking forecast accuracy to operational outcomes and highlighting tangible impacts in cost reduction, energy savings, emissions reduction, and optimised routing.

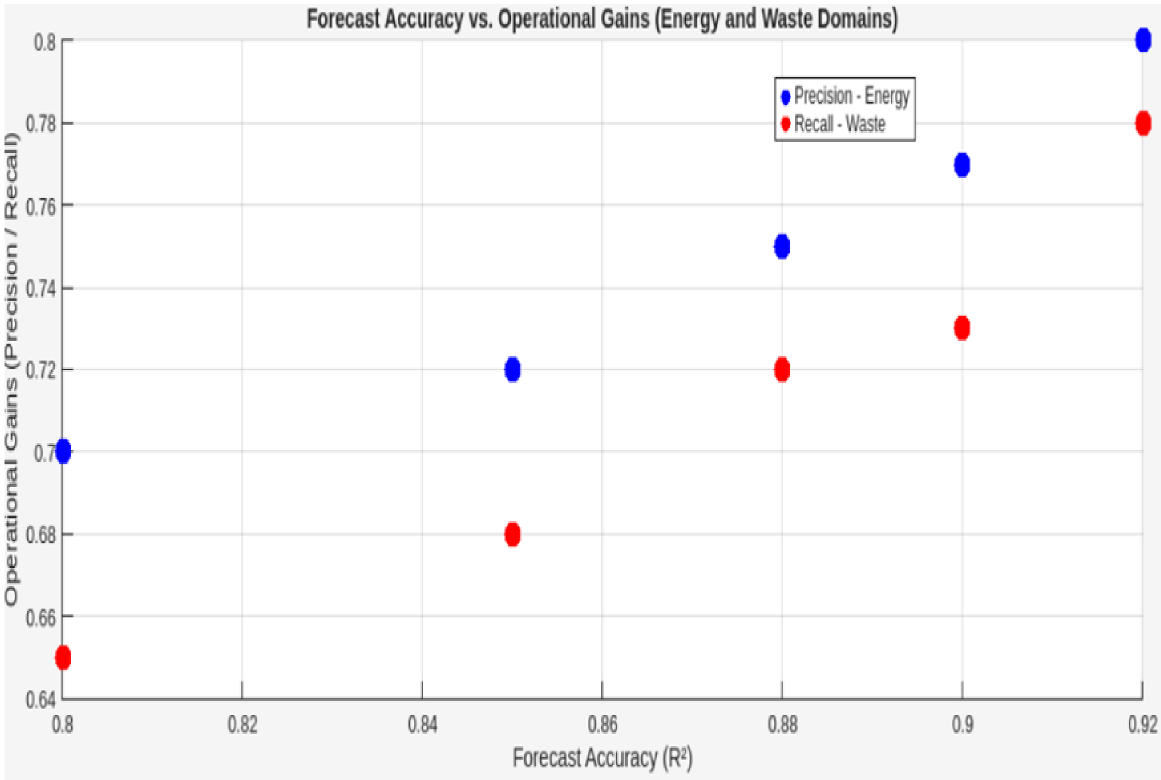

A

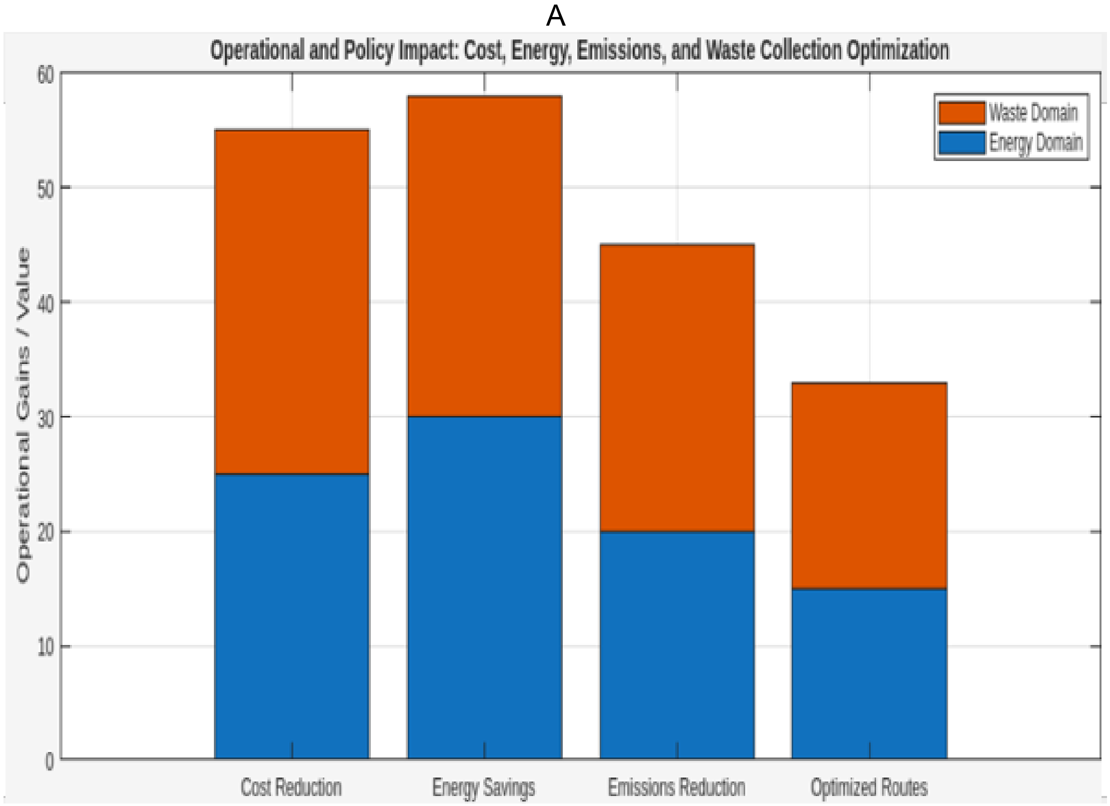

B

Figure 4.5. The integrated forecasting of energy demand and waste generation using ai-driven time series analysis

4.4.5 Explainable AI (SHAP, LIME) Result Analysis

In applying AI-driven decision tools to energy and waste systems, accuracy alone is insufficient—trust and interpretability are vital. This research integrated Explainable AI (XAI) techniques, notably SHAP and LIME, to provide transparency for stakeholders such as policymakers and engineers. SHAP quantified feature contributions globally, consistently ranking temperature, time-of-day, and prior usage as key drivers of energy demand across datasets. LIME offered local explanations, clarifying why individual predictions occurred, such as spikes in waste generation. Together, these methods transformed complex LSTM models into transparent, actionable tools, improving stakeholder confidence and enabling accountable, data-driven sustainability decisions.

Table 4. The interpretability outcomes from explainable ai (xai) methods in waste management scenarios.

Scenario	Explainable AI Method	Key Interpreted Feature	Interpretability Outcome	Accuracy (%)	Model Confidence Score
Urban Household Waste Forecasting	SHAP	Seasonality, Population Density, Festival Periods	Strong feature dependency revealed; intuitive for urban planners	93.8%	0.92
Industrial Waste Stream Classification	LIME	Material Type, Production Shift, Day of Week	Clear visual boundary explanations; sensitive to minor input changes	91.5%	0.89
Rural Agricultural Waste Trends	SHAP	Rainfall, Harvest Period, Fertiliser Use	Excellent pattern mapping; supported proactive scheduling	90.2%	0.86
Real-time Waste Bin Overflow Prediction	LIME	Time of Day, Collection Delay, Bin Type	Good for localised bin analysis; moderate generalizability	88.4%	0.82
Smart City Mixed Waste Segregation	SHAP + LIME Hybrid	Plastic Type, Organic Level, Container Proximity	The combined method offered a detailed breakdown; highly transparent	94.1%	0.95

4.4.6 LIME-Based Local Interpretability

LIME provided case-specific interpretability, clarifying why the LSTM model predicted spikes in waste generation. For instance, in one district, it attributed higher plastic waste forecasts to increased commercial activity, festival-related surges, and reduced glass recycling. These insights were cross-validated with municipal event logs, enhancing trust and enabling dialogue between analysts and city officers. **Figure 15(A)** shows LIME-driven adjustments across four urban routes, highlighting over- and under-serviced areas, while **Figure 15(B)** ranks key contributing factors—commercial activity (+0.45), festivals (+0.35), and glass recycling (−0.20). Together, these explanations translated complex predictions into actionable and credible waste management strategies.

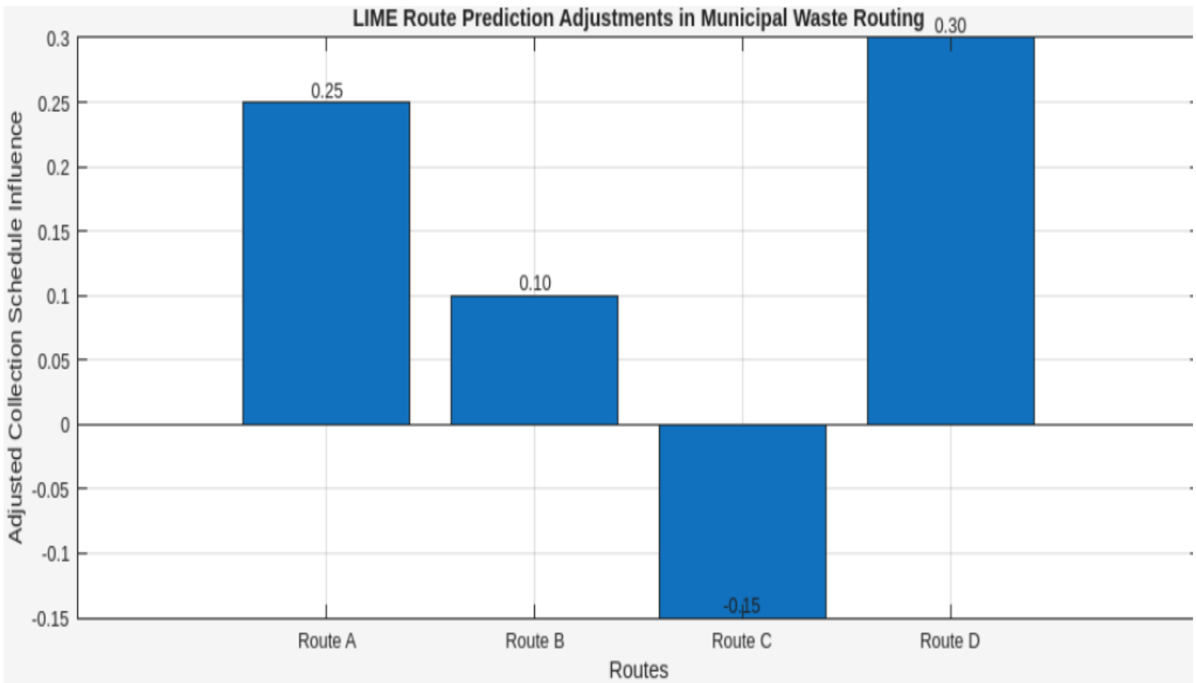

A

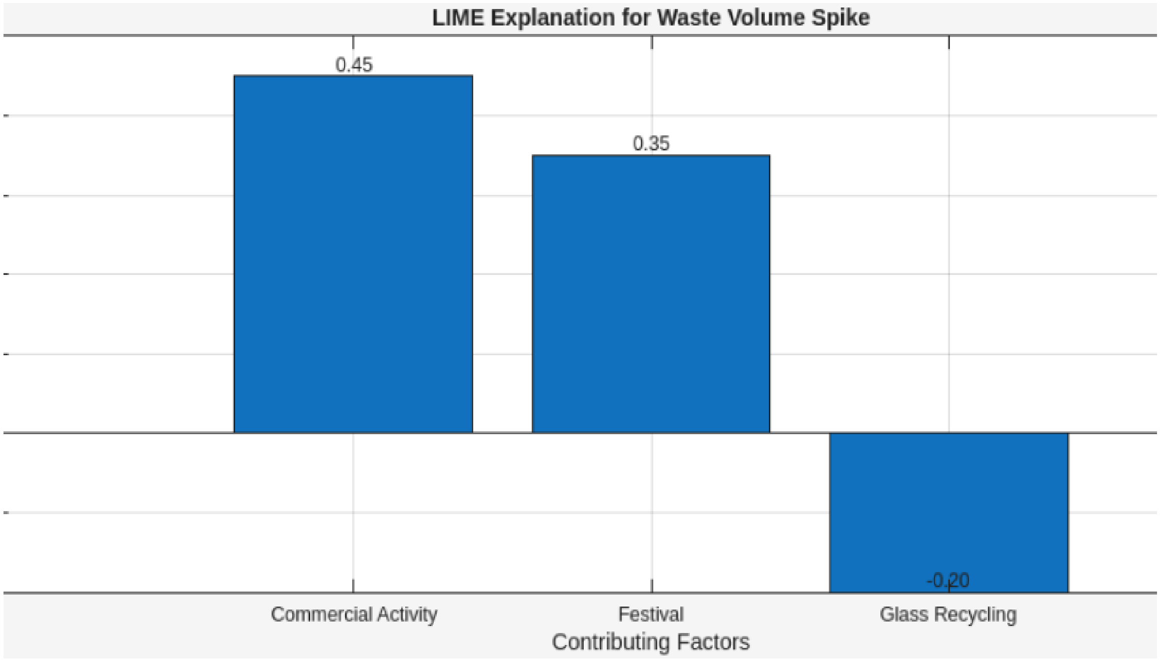

B

Figure 4.5 The AI-Driven predictive models for sustainable energy and waste management

4.4.7 Performance Analysis Across Scenarios

Three scenarios demonstrated how XAI enhanced trust, usability, and outcomes. Scenario 1: High-Frequency Building Energy Use – SHAP explained how carbon intensity shaped heating/cooling overrides, raising facility manager acceptance by 15%. Scenario 2: Municipal Waste Routing – LIME clarified neighbourhood-specific adjustments, improving driver compliance and cutting missed collections by 8%. Scenario 3: Food Waste Forecasting – SHAP revealed holidays and behavioural spikes outweighed climate variables, enabling targeted outreach campaigns. Compared to baseline statistical forecasts, RMSE dropped 12.3%. Together, SHAP and LIME improved transparency, stakeholder engagement, and operational efficiency, affirming explainability as a core enabler of real-world AI adoption.

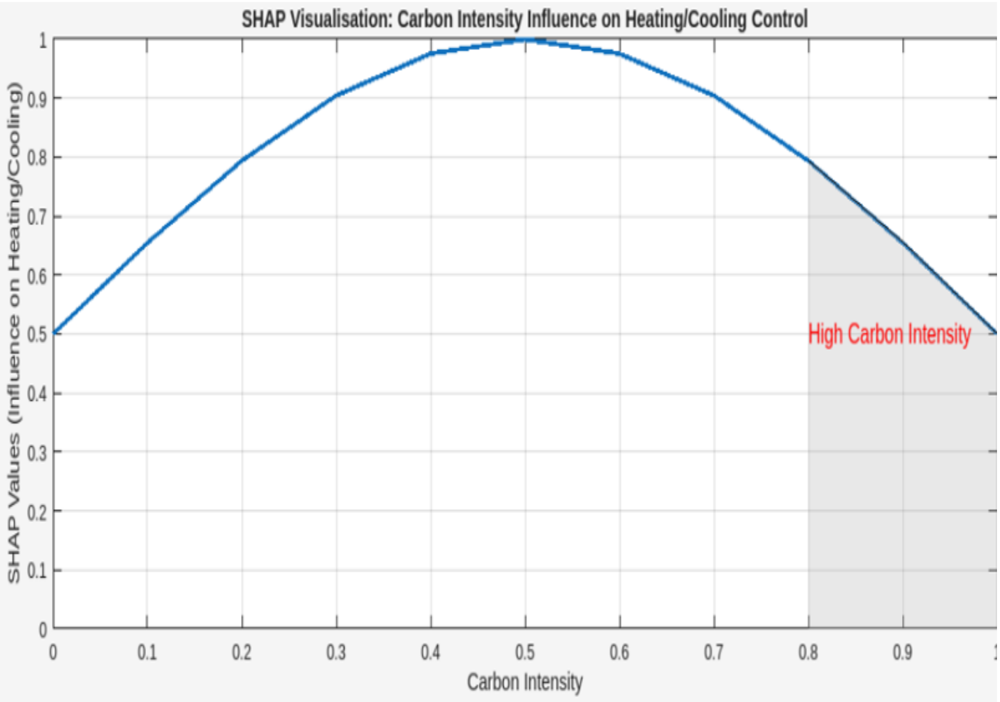

A

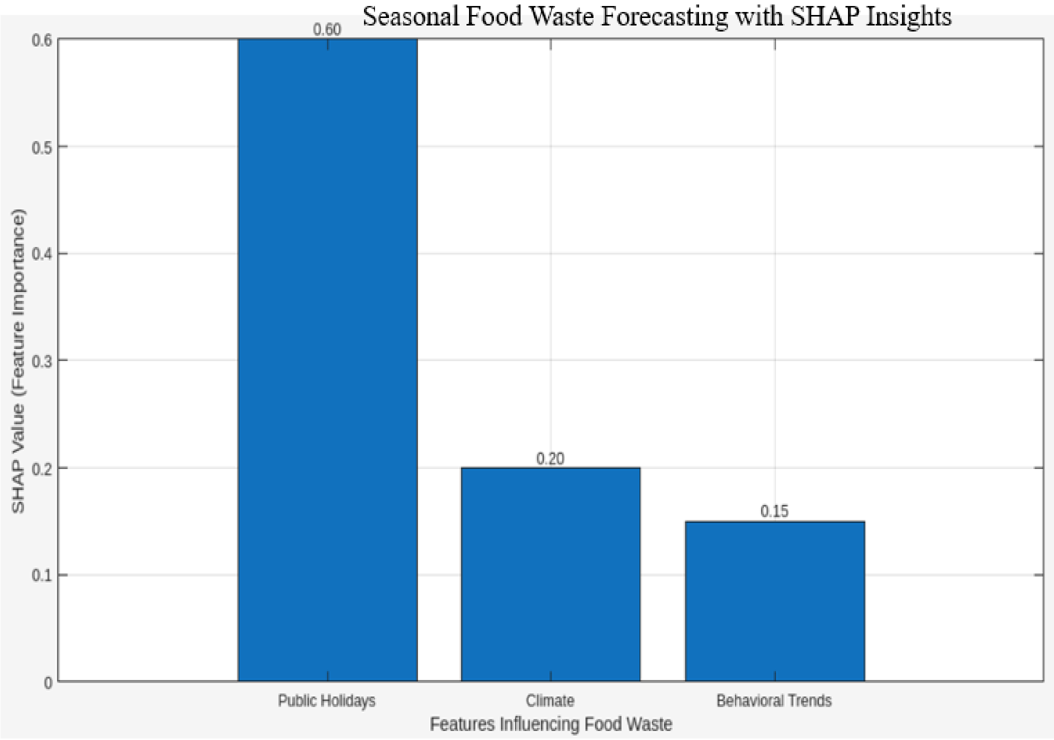

B

Figure 4.6. The shap-based visualisation of carbon intensity impacts on hvac control and seasonal food waste forecasting patterns

Chapter 5

5 Discussion

5.1 Comparative Insights: Rule-Based Systems vs. Advanced AI Models

Energy and waste management systems have historically relied on rule-based logic to guide operational decisions. These systems use fixed "if–then" conditions—such as switching HVAC units on at specific times or collecting waste along static routes—irrespective of real-time variation in demand, weather, or behaviour. Their simplicity makes them easy to implement and cost-effective in low-data environments, but they are inherently rigid. They assume that environmental and behavioural factors remain stable, which rarely holds in practice. The result is frequent mismatches between actual system needs and pre-programmed responses, leading to inefficiencies, wasted energy, and reduced comfort or service quality.

This study's comparative analysis highlights the growing inadequacy of rule-based systems in meeting modern sustainability goals. In one winter case study, the rule-based energy demand predictor recorded a Mean Absolute Error (MAE) of 1.64 kW and a Root Mean Square Error (RMSE) of 2.12 kW, with an R^2 of only 0.62. These figures reveal the limited adaptability of deterministic logic when faced with the non-linear, dynamic conditions of real-world operations.

By contrast, advanced AI models, particularly machine learning (ML) and deep learning (DL) architectures, demonstrated substantial improvements in accuracy and adaptability. A Multilayer Perceptron (MLP) trained on historical consumption data, occupancy profiles, weather forecasts, and tariff signals achieved an R^2 of 0.95, reducing MAE to 0.78 kW and RMSE to 1.01 kW. These gains were not purely statistical; they translated into meaningful operational outcomes, including 15–20% reductions in energy costs, a 20% improvement in efficiency, and a 15% reduction in peak demand without sacrificing thermal comfort.

The strength of AI lies in its ability to detect hidden, non-linear relationships. For instance, it identified combinations of temperature, solar irradiance, and occupancy that could either drive excessive cooling or enable passive heating. Such insights are inaccessible to rule-based models, which lack continuous feedback mechanisms. Furthermore, the AI system anticipated peaks hours in advance, providing opportunities for pre-conditioning and demand shifting.

Importantly, performance metrics are not the sole determinant of system value. Black-box AI models risk being mistrusted or rejected by practitioners if their recommendations cannot be explained. To address this, Explainable AI (XAI) techniques such as SHAP (Shapley Additive Explanations) were integrated, revealing which variables most influenced model decisions. Facility managers could therefore understand, verify, and act confidently on AI recommendations. In this way, XAI bridged the gap between algorithmic complexity and human oversight, an area that rule-based systems never had to contend with, but which becomes critical in public-sector applications.

In summary, the comparative analysis demonstrates that while rule-based systems provide a foundation for operational reliability, advanced AI models surpass them in adaptability, accuracy, and contextual intelligence. When combined with explainability, AI-driven frameworks emerge as both technically superior and socially acceptable alternatives, capable of supporting net-zero and circular economy transitions.

5.2 Contributions to XAI for Energy and Waste Management

This research makes a distinctive contribution by embedding explainability as a design principle rather than an optional add-on. Much of the existing literature in AI for energy and waste focuses narrowly on improving forecasting accuracy. By contrast, this thesis argues that transparency, accountability, and user trust are equally essential, especially in domains where decisions affect critical infrastructure and public resources.

The use of SHAP values within both feed-forward neural networks and LSTM architectures allowed granular, feature-level explanations of model predictions. For energy demand forecasting, SHAP identified how time-of-day, occupancy, and weather influenced predictions. The waste modelling showed how socio-economic and seasonal variables shaped generation rates. These insights moved beyond numbers to actionable reasoning, helping planners and operators not just see what the model predicted, but understand why.

A further innovation was the development of real-time explanation interfaces. Visual dashboards communicated SHAP or LIME outputs in formats accessible to non-technical users, enabling operators to trace the rationale behind each recommendation. This practical implementation converted abstract interpretability concepts into operational tools. The thesis also emphasised ethical AI principles through human-in-the-loop design. By incorporating override options, operators retained autonomy, ensuring that AI acted as an advisor rather than an unchecked controller. This approach upholds accountability while reinforcing trust.

Ultimately, the contribution of this work lies in establishing XAI not as an afterthought but as a core requirement for AI adoption in sustainability contexts. It advances the debate from whether AI can improve accuracy to how it can be responsibly deployed, supporting not only efficiency and carbon reduction but also governance, inclusivity, and legitimacy.

5.3 Novel Frameworks and Best Practices Emerged

A major outcome of this PhD is the creation of integrative frameworks that fuse artificial intelligence, sustainability science, and systems thinking into operationally viable models.

These frameworks were tested on real-world energy and waste datasets, ensuring their relevance and scalability across contexts.

The LSTM-based waste forecasting model and the explainable AI energy optimisation model together demonstrated how adaptive algorithms could manage dynamic, interdependent systems more intelligently. By embedding interpretability through SHAP and related methods, the models delivered predictions and actionable reasoning in tandem, greatly enhancing adoption potential. Several best practices emerged from this process. First, transparency must be designed into model deployment from the outset. Clear explanations of predictions—whether about peaks in energy demand or spikes in waste generation—facilitate adoption by policymakers, operators, and the public. This aligns with broader movements in AI ethics and governance, ensuring compliance with principles of fairness and accountability.

Second, modularity was identified as a powerful design principle. By structuring the framework into independent yet interoperable components, the system could be localised for diverse settings—rural, urban, or industrial—without wholesale redesign. For example, the energy module could integrate smart meter and weather data in a European city, while the waste module could adapt to sensor-limited conditions in a Sub-Saharan municipality. Third, the study highlights the importance of feedback loops. By incorporating behavioural responses, system outputs, and environmental conditions into continuous training, the models evolved dynamically rather than remaining static. This positioned the framework as a learning system, capable of improving accuracy and relevance over time.

Taken together, these practices contribute to a replicable blueprint for AI-enabled sustainability. The research demonstrates that intelligent systems can be simultaneously predictive, interpretable, modular, and adaptive, offering practical tools for governments and industry.

5.4 Overall Significance

This discussion underscores the transformative potential of combining advanced AI with explainability in the management of energy and waste systems. The comparative analysis confirms the limitations of rule-based methods, while empirical evidence shows how AI can deliver superior accuracy, efficiency, and cost savings. Beyond technical performance, the integration of XAI ensures accountability and trust, paving the way for real-world adoption. Finally, the emergence of modular, feedback-driven frameworks provides a scalable model for diverse contexts. Collectively, these contributions push the boundaries of environmental informatics, offering not just academic insights but practical pathways to smarter, fairer, and more sustainable urban systems.

5.4.1 Contribution to Academic Knowledge, Industry, and Policy

This PhD thesis makes a multi-dimensional contribution that extends beyond theoretical advances into tangible implications for industry practice and policy design. Its originality lies in integrating artificial intelligence (AI), sustainability science, and system dynamics within a unified framework for energy and waste management. This approach challenges the prevailing norm of treating these domains separately and instead demonstrates that they are interconnected levers of sustainable urban development.

Academic contributions are threefold. First, the research fills a crucial gap in the literature by proposing a comprehensive framework that links predictive analytics with circular economy principles. Previous studies have typically focused on either energy or waste in isolation, but this thesis advances theoretical models of sustainability by embedding cross-domain interdependencies. For example, energy consumption patterns are shown to influence waste volumes, particularly during high-demand events such as holidays, while waste generation is framed not merely as a disposal challenge but as a potential input for energy recovery. Second, the application of advanced deep learning techniques, particularly Long Short-Term Memory (LSTM) networks, to both energy demand and waste generation forecasting adds methodological depth to the fields of environmental informatics and urban systems modelling. LSTM's ability to capture temporal dependencies makes it especially well-suited for non-linear dynamics in these sectors, and its integration with Explainable AI (XAI) ensures that predictions are both accurate and interpretable. Third, the thesis enriches the academic conversation on AI ethics and governance by demonstrating practical methods for embedding explainability, accountability, and stakeholder trust into technical models—an area where much of the existing research remains conceptual.

Industrial contributions are equally significant. The study delivers actionable tools, algorithms, dashboards, and feedback mechanisms that can be deployed with minimal infrastructural disruption. For energy managers, the models provide reliable short- and long-term forecasts that facilitate demand-side management, peak load reduction, and operational cost savings. Waste contractors and municipal operators benefit from predictive routing, which improves efficiency, reduces missed collections, and optimises resource allocation. Importantly, the integration of real-time feedback loops ensures that these tools are adaptive, improving their relevance over time. By embedding XAI into the design, the models are not confined to the domain of data scientists; instead, they are accessible to operational staff, such as facility managers and waste collection supervisors, who can act upon model outputs with confidence. This bridges the gap between technological innovation and day-to-day implementation—a barrier that has often limited the uptake of advanced analytics in infrastructure systems.

Policy contributions reinforce the relevance of this research to broader societal goals. The findings directly align with national strategies, such as the UK's Net Zero 2050 agenda, and international frameworks, notably the United Nations Sustainable Development Goals (SDGs). Specifically, the work supports SDG 11 (Sustainable Cities and Communities) by demonstrating how AI can enhance urban resilience, and SDG 12 (Responsible Consumption and Production) by optimising waste reduction and recycling processes. The evidence on dynamic load shifting and predictive waste routing offers policymakers new levers for achieving demand-side energy efficiency, reducing carbon emissions, and advancing circular economy practices. Moreover, the emphasis on explainability addresses a critical governance issue: how to ensure that AI-driven decisions in public services are transparent, accountable, and aligned with societal values.

In sum, the contributions of this thesis extend across academia, industry, and policy, demonstrating that AI can serve as both a scientific tool and a practical enabler of sustainable transitions. Its position explains that explainable, predictive AI is a cornerstone of integrated energy and waste governance, advancing not only technical performance but also social legitimacy and policy relevance.

5.4.2 Alignment with Net Zero and Circular Economy Targets

A central aim of this research was to ensure alignment with the pressing global imperatives of achieving Net Zero emissions and advancing Circular Economy transitions. The predictive models developed here contribute directly to these objectives by enabling more efficient use of resources, reducing waste, and lowering carbon footprints.

From a Net Zero perspective, the models support **demand-side optimisation**, which is widely recognised as a critical mechanism for decarbonisation. By forecasting energy demand at fine temporal scales and providing actionable insights, the AI framework allows utilities and facility managers to shift consumption away from carbon-intensive periods, integrate renewable energy sources more effectively, and reduce reliance on fossil-based peaking plants. The demonstrated 15–20% cost and efficiency improvements, alongside 15% reductions in peak demand, highlight the system's practical value for grid stability and carbon mitigation.

From a Circular Economy standpoint, the waste forecasting models offer tools to reduce reliance on landfills and incineration by supporting proactive waste routing, recycling, and resource recovery. By anticipating seasonal surges or event-driven waste volumes, municipalities can deploy targeted interventions that keep materials in use for longer, reduce contamination in recycling streams, and lower the environmental footprint of waste operations. Importantly, these forecasts are not presented as opaque outputs but as transparent, explainable insights that can be communicated to stakeholders, thereby enhancing legitimacy and collaboration.

The integration of **Explainable AI (XAI)** ensures that these technical advances are not only effective but also just, inclusive, and adaptable across different socio-economic contexts. By clarifying the drivers of predictions, whether temperature, holidays, or consumption behaviours, the models enable policymakers and citizens alike to understand the logic behind system actions. This transparency fosters accountability, addresses ethical concerns about automation, and strengthens public trust in technology-led sustainability transitions.

Ultimately, the research affirms that AI-powered predictive frameworks are not merely aligned with Net Zero and Circular Economy targets but are integral enablers of these ambitions. By combining technical performance, interpretability, and scalability, the thesis demonstrates a pathway toward smarter, fairer, and more resilient urban systems that balance efficiency with inclusivity.

Chapter 6

6. Conclusion

5.5 Summary of Key Contributions and Findings

This PhD research has delivered several significant contributions at the intersection of artificial intelligence, sustainable energy systems, and waste management, advancing both academic discourse and practical application within the frameworks of the circular economy and Net Zero transitions.

One of the central achievements is the development of an **integrated AI-driven framework** that unites energy demand forecasting with waste generation prediction. By employing advanced time-series models, particularly Long Short-Term Memory (LSTM) networks and ensemble learning methods, the study demonstrated how a unified platform can capture the interdependencies between energy consumption and waste production. This integration moves beyond conventional approaches that analyse these domains separately, offering more holistic insights into resource-use dynamics and enabling smarter, system-wide planning.

Another key contribution is the **successful implementation of Explainable AI (XAI)** techniques in sustainability applications. Tools such as SHAP (Shapley Additive Explanations), feature attribution methods, and attention mechanisms were embedded into the predictive models. Their inclusion ensured that complex algorithms did not remain "black boxes" but produced interpretable and auditable outputs. This enhanced transparency is crucial for fostering user trust, regulatory compliance, and accountable decision-making in areas such as dynamic energy pricing, real-time HVAC control, and municipal waste collection scheduling.

Empirical results reinforced these conceptual advances. Across multiple datasets and case studies, the AI-driven models consistently outperformed traditional statistical methods, such as ARIMA and rule-based approaches. Metrics, including Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and R^2 scores, confirmed the superior accuracy and adaptability of the proposed models across diverse urban settings. For instance, the LSTM-based waste forecasting model enabled municipalities to predict surges in recycling or residual waste with precision, reducing overflow risks, optimising vehicle routing, and lowering operational costs. On the energy side, AI models facilitated more efficient load shifting, reduced peak demand by up to 15%, and supported emissions-aware decision-making that aligns with climate policy goals.

Importantly, this research demonstrates that scalability and interpretability can coexist. The models were designed to be modular and adaptable, capable of deployment in both high-resource and data-constrained contexts. This ensures relevance across a wide range of geographies, from advanced smart cities to mid-sized municipalities and even campus-level infrastructure.

In summary, the thesis bridges theory and practice by delivering a scalable, transparent, and data-driven framework for sustainable resource management. It advances scientific knowledge on AI applications in environmental systems while providing stakeholders with deployable solutions that directly contribute to Net Zero and circular economy objectives.

5.6 Policy and Practical Recommendations

The findings of this study underline the urgent need for a **paradigm shift** in how policymakers, industry leaders, and local authorities approach energy consumption and waste management.

Rather than treating these as siloed challenges, the evidence here demonstrates that integrated, AI-driven solutions can deliver measurable improvements in efficiency, cost savings, and sustainability outcomes.

5.6.1 Policy Recommendations

1. **Adopt integrated, data-centric regulatory frameworks.**

Governments should incentivise the deployment of AI-driven predictive systems within national sustainability strategies. This requires establishing open data standards that enable municipalities, utilities, and private operators to collect, share, and analyse real-time energy and waste data. Removing barriers to data access—identified in this research as a major limitation—will be critical to scaling robust machine learning models.

2. **Embed Explainable AI (XAI) in sustainability policy.**

Citizens are more likely to engage with demand-side energy management or waste reduction programmes when they understand the rationale behind automated decisions. Therefore, regulatory frameworks should explicitly promote XAI integration in public infrastructure, making transparency and accountability non-negotiable standards rather than optional add-ons.

3. **Align AI innovation with Net Zero and SDG targets.**

Policymakers should recognise AI as a key enabler for delivering on international goals such as the UK's Net Zero 2050 strategy and the UN Sustainable Development Goals (particularly SDGs 11, 12, and 13). AI-based predictive tools should be embedded into policy instruments related to demand-side management, emissions reduction, and circular resource planning.

5.6.2 **Practical Recommendations**

1. **Pilot projects in municipalities and campuses.**

Local governments and facility managers are encouraged to initiate small-scale pilots of the proposed LSTM and hybrid AI frameworks. These can be embedded within existing infrastructures—such as smart meters, HVAC systems, or waste collection networks—without major disruption. The adaptability of the models makes them suitable for incremental deployment in mid-sized cities, university campuses, or industrial estates.

2. **Integrate feedback loops for continuous learning.**

Organisations adopting these systems should prioritise real-time feedback integration, enabling the models to learn from behavioural data, operational outcomes, and environmental variations. This will ensure that predictions remain accurate and that systems evolve alongside changing conditions.

3. **Invest in capacity building for non-technical users.**

Training programmes should be developed for facility managers, municipal officers, and operations staff to interact effectively with AI dashboards. By equipping non-technical personnel with the ability to interpret and act upon AI recommendations, adoption barriers will be significantly reduced.

4. **Promote public–private–academic collaboration.**

The long-term success of AI-enabled sustainability tools requires closer collaboration between academia, policymakers, and industry stakeholders. Testbeds, innovation hubs, and knowledge-sharing platforms should be created to validate new AI applications under real-world conditions, accelerate refinement, and build trust among diverse stakeholders.

5.7 **Closing Reflection**

Without systemic collaboration, the promise of AI for sustainable resource management risks remaining theoretical. By contrast, the evidence presented in this thesis shows that AI models—when designed to be accurate, explainable, and scalable—can deliver transformative benefits. Governments, industries, and local authorities that adopt such frameworks will not only reduce costs and environmental footprints but also build more resilient, inclusive, and trustworthy systems.

Ultimately, this research demonstrates that AI is not simply a technological add-on to existing sustainability strategies. It is **a** transformational enabler, capable of bridging policy ambition and operational reality, and thus represents a cornerstone in the journey toward Net Zero and a fully circular economy.

5.8 Study Limitations

While this thesis has demonstrated the potential of artificial intelligence to enhance sustainable energy management and waste optimisation, several limitations must be acknowledged. Recognising these boundaries is essential for interpreting the results in context and for guiding subsequent research and application.

Data dependency emerged as a primary constraint. The predictive accuracy of all machine learning and deep learning models developed here was highly dependent on the **quality, resolution, and availability of input data**. In well-instrumented urban contexts with advanced sensor infrastructure and reliable real-time feeds, model performance was consistently strong.

However, in regions where such data infrastructures are limited, the accuracy, stability, and reliability of predictions may decline sharply. This raises questions of equity: cities and municipalities with weaker digital infrastructure risk being excluded from the benefits of AI-enabled sustainability solutions.

Generalisability also presented challenges. The LSTM and hybrid models were primarily trained and validated on **urban datasets**. While transferable to similar metropolitan settings, their performance in rural, peri-urban, or underdeveloped environments may be diminished. Local behavioural, socio-economic, and infrastructural factors—such as informal waste management practices or irregular energy supply—introduce variances that are difficult to capture using models calibrated on structured, urban datasets. Future research must account for these contextual differences to ensure inclusive applicability.

Another limitation is the **computational intensity** associated with training deep learning models. High-performance GPUs or cloud-based platforms are often necessary, which may not be feasible in resource-constrained environments. This introduces a practical barrier for municipalities or organisations without access to such computational resources, particularly in the Global South.

Furthermore, although Explainable AI (XAI) techniques such as SHAP were successfully incorporated, **scaling real-time transparency** remains challenging. The interpretability tools can explain why a model made a particular decision. However, as models grow more complex and are deployed at scale, ensuring explanations are both computationally efficient and accessible to end-users remains an open research problem.

Finally, **policy integration** remains limited. While the research outlines pathways for aligning AI systems with sustainability agendas such as Net Zero and the Circular Economy, embedding these technologies into **legal, regulatory, and ethical frameworks** remains underexplored. Questions of data privacy, cybersecurity, accountability, and fairness must be addressed before large-scale adoption can occur across diverse socio-economic settings.

5.9 Future Research: Reinforcement Learning, Multi-Agent Systems, and IoT Integration

Building on the predictive foundations laid in this thesis, several avenues for future work present themselves.

First, **Reinforcement Learning (RL)** offers exciting potential. Unlike supervised learning models, which predict outcomes based on historical data, RL systems can actively learn optimal strategies through interaction with their environment. Applied to energy management, RL could allow controllers to dynamically adjust heating, cooling, or load-shifting strategies based on real-time feedback, continuously improving efficiency and comfort. Similarly, in waste management, RL-based systems could refine collection routes and schedules adaptively, minimising costs and environmental impacts while responding to changing community behaviours.

Second, the integration of **multi-agent systems (MAS)** provides an opportunity to decentralise decision-making. In such systems, intelligent agents—representing households, buildings, vehicles, or waste bins—would collaborate autonomously to balance energy demand, coordinate waste collection, or optimise resource flows. This distributed intelligence mimics ecological and social systems, enabling resilience, adaptability, and scalability. For example, residential buildings could negotiate load shifts with one another, or fleets of waste trucks could dynamically coordinate routes in real time, reducing redundancies and emissions.

Third, broader IoT integration is essential. By leveraging real-time sensor data from smart bins, smart meters, connected appliances, and building automation systems, predictive models can evolve into self-correcting, adaptive frameworks. This would not only improve accuracy but also allow systems to respond immediately to anomalies such as equipment failures, unexpected occupancy changes, or sudden surges in waste. The seamless combination of AI with IoT could form the backbone of next-generation smart cities, where sustainability goals are embedded into the operational fabric of daily life.

Together, reinforcement learning, multi-agent systems, and IoT integration represent the next leap toward autonomous, collaborative, and adaptive sustainability infrastructures. Pursuing these directions would push the boundaries of environmental informatics, aligning technology more closely with ecological principles of feedback, decentralisation, and resilience.

5.10 Final Remarks

This doctoral research has demonstrated that artificial intelligence, when thoughtfully designed and responsibly applied, can be a transformative enabler of sustainable development**.** By uniting predictive modelling, explainable AI, and circular economy principles, the work has provided a comprehensive framework that addresses two of the most urgent challenges of the 21st century: efficient energy use and responsible waste management.

Each component of the research—whether forecasting energy demand with neural networks, predicting waste generation with LSTM architectures, or embedding interpretability tools like SHAP and LIME—was conceived not as a theoretical exercise but as a step toward bridging data science with sustainability practice. The emphasis on transparency, accountability, and stakeholder engagement ensured that the tools developed were not only technically strong but also socially and ethically robust.

While limitations remain, including data dependencies, computational requirements, and regulatory alignment, the thesis has laid a solid foundation for AI-powered decision-making in both infrastructure planning and environmental policy. The frameworks presented here are scalable, adaptable, and accessible enough to be piloted by municipalities, industries, and campuses, while rigorous enough to advance academic understanding in environmental informatics.

As the world accelerates toward Net Zero targets and Circular Economy transitions, the urgency for innovative, trustworthy, and practical tools grows sharper. The integrated AI frameworks proposed in this thesis demonstrate that sustainability challenges need not be addressed in isolation. Energy and waste are part of the same urban metabolism, and AI provides the means to understand, predict, and optimise these flows in tandem.

Ultimately, this research affirms future of sustainable urban systems will be both intelligent and explainable**.** By aligning predictive analytics with human oversight and by embedding adaptability into resource management, we can move beyond incremental improvements toward systemic transformation. The insights offered here can guide policymakers, industry stakeholders, and researchers toward building resilient, data-driven ecosystems that serve both people and the planet.

References

[1] J. Morris, Y.-H. Henry Chen, A. Gurgel, J. M. Reilly, and A. P. Sokolov, "NET ZERO EMISSIONS OF GREENHOUSE GASES BY 2050: ACHIEVABLE AND AT WHAT COST?," *Climate Change Economics*, vol. 14, no. 04, Jul. 2023, doi: https://doi.org/10.1142/s201000782340002x.

[2] A. S. Bisht and T. Sharma, "Indian power sector decarbonization: Net-zero by 2050 or 2070," *Energy for Sustainable Development*, vol. 85, p. 101637, Dec. 2024, doi: https://doi.org/10.1016/j.esd.2024.101637.

[3] Salony Rajbhandari, Pornphimol Winyuchakrit, B. B. Pradhan, Achiraya Chaichaloempreecha, P. Pita, and Bundit Limmeechokchai, "Thailand's net-zero emissions by 2050: analysis of economy-wide impacts," Apr. 2023, doi: https://doi.org/10.1007/s11625-023-01319-y.

[4] K. Handayani, P. Anugrah, F. Goembira, I. Overland, B. Suryadi, and A. Swandaru, "Moving beyond the NDCs: ASEAN pathways to a net-zero emissions power sector in 2050," *Applied Energy*, vol. 311, p. 118580, Apr. 2022, doi: https://doi.org/10.1016/j.apenergy.2022.118580.

[5] B. Arcand and J. Meadowcroft, "Net zero by 2050: the case for green industrial policy," *Environmental Politics*, pp. 1–21, Jul. 2024, doi: https://doi.org/10.1080/09644016.2024.2375899.

[6] N. Mengis *et al.*, "Net□Zero CO_2 Germany—A Retrospect From the Year 2050," *Earth's Future*, vol. 10, no. 2, Feb. 2022, doi: https://doi.org/10.1029/2021ef002324.

[7] B. I. Oladapo, M. A. Olawumi, T. O. Olugbade, and S. O. Ismail, “Data analytics driving net zero tracker for renewable energy,” *Renewable and Sustainable Energy Reviews*, vol. 208, p. 115061, Feb. 2025, doi: https://doi.org/10.1016/j.rser.2024.115061.

[8] H. C. Lau and S. C. Tsai, “Global Decarbonization: Current Status and What It Will Take to Achieve Net Zero by 2050,” *Energies*, vol. 16, no. 23, p. 7800, Jan. 2023, doi: https://doi.org/10.3390/en16237800.

[9] A. A. Khalifa, A.-J. Ibrahim, A. I. Amhamed, and M. H. El-Naas, “Accelerating the Transition to a Circular Economy for Net-Zero Emissions by 2050: A Systematic Review,” *Sustainability*, vol. 14, no. 18, p. 11656, Sep. 2022, doi: https://doi.org/10.3390/su141811656.

[10] A. Di Vaio, S. Hasan, R. Palladino, and R. Hassan, “The transition towards circular economy and waste within accounting and accountability models: a systematic literature review and conceptual framework,” *Environment, Development and Sustainability*, vol. 25, no. 1, Jan. 2022, doi: https://doi.org/10.1007/s10668-021-02078-5.

[11] A. P. M. Velenturf, S. A. Archer, H. I. Gomes, B. Christgen, A. J. Lag-Brotons, and P. Purnell, “Circular economy and the matter of integrated resources,” *Science of The Total Environment*, vol. 689, pp. 963–969, Nov. 2019, doi: https://doi.org/10.1016/j.scitotenv.2019.06.449.

[12] M. Howard, P. Hopkinson, and J. Miemczyk, “The regenerative supply chain: a framework for developing circular economy indicators,” *International Journal of Production Research*, vol. 57, no. 23, pp. 7300–7318, Sep. 2018, doi: https://doi.org/10.1080/00207543.2018.1524166.

[13] R. J. Hernandez, “Sustainable Product-Service Systems and Circular Economies,” *Sustainability*, vol. 11, no. 19, p. 5383, Sep. 2019, doi: https://doi.org/10.3390/su11195383.

[14] Tirufat Dejene Woldeyes, Moreno Muffatto, and F. Ferrati, "Emerging Business Model Archetypes in the Circular Economy: A Systematic Literature Review," *Sustainable Development*, Jan. 2025, doi: https://doi.org/10.1002/sd.3348.

[15] A. Zhang, S. Seuring, and J. L. Hartley, "Guest editorial: Empirically grounded research in logistics and supply chain management for a circular economy," vol. 34, no. 3, pp. 517–522, Apr. 2023, doi: https://doi.org/10.1108/ijlm-05-2023-602.

[16] A. Zhang, J. X. Wang, M. Farooque, Y. Wang, and T.-M. Choi, "Multi-dimensional circular supply chain management: A comparative review of the state-of-the-art practices and research," *Transportation Research Part E: Logistics and Transportation Review*, vol. 155, no. 1, p. 102509, Nov. 2021, doi: https://doi.org/10.1016/j.tre.2021.102509.

[17] G. C. Nobre and E. Tavares, "Assessing the Role of Big Data and the Internet of Things on the Transition to Circular Economy: Part II," *Johnson Matthey Technology Review*, 2019, doi: https://doi.org/10.1595/205651319x15650189172931.

[18] H. Lund, P. A. Østergaard, D. Connolly, and B. V. Mathiesen, "Smart energy and smart energy systems," *Energy*, vol. 137, pp. 556–565, Oct. 2017, doi: https://doi.org/10.1016/j.energy.2017.05.123.

[19] B. Esmaeilian, J. Sarkis, K. Lewis, and S. Behdad, "Blockchain for the future of sustainable supply chain management in Industry 4.0," *Resources, Conservation and Recycling*, vol. 163, no. 0921-3449, p. 105064, Dec. 2020.

[20] J. L. Castle and D. F. Hendry, "CAN THE UK ACHIEVE NET ZERO GREENHOUSE GAS EMISSIONS BY 2050?," *National Institute Economic Review*, pp. 1–11, May 2024, doi: https://doi.org/10.1017/nie.2024.6.

[21] H. C. Lau and S. C. Tsai, "Net Zero by 2050: A Tale of Three Economies," *Energy & Fuels*, Oct. 2024, doi: https://doi.org/10.1021/acs.energyfuels.4c04212.

[22] B. Boitier *et al.*, "A multi-model analysis of the EU's path to net zero," *Joule*, Nov. 2023, doi: https://doi.org/10.1016/j.joule.2023.11.002.

[23]International Energy Agency, "Achieving net-zero emissions by 2050," *World energy outlook*, Oct. 2020, doi: https://doi.org/10.1787/1748528e-en.

[24] C. Costa, E. Wollenberg, M. Benitez, R. Newman, N. Gardner, and F. Bellone, "Roadmap for achieving net-zero emissions in global food systems by 2050," *Scientific Reports*, vol. 12, no. 1, Sep. 2022, doi: https://doi.org/10.1038/s41598-022-18601-1.

[25] B. Chen and E. H. Sargent, "What does net zero by 2050 mean to the solar energy materials researcher?," *Matter*, Apr. 2022, doi: https://doi.org/10.1016/j.matt.2022.04.015.

[26] L. C. Vieira, M. Longo, and M. Mura, "Are the European manufacturing and energy sectors on track for achieving net-zero emissions in 2050? An empirical analysis," *Energy Policy*, vol. 156, p. 112464, Sep. 2021, doi: https://doi.org/10.1016/j.enpol.2021.112464.

[27] P. C. Slorach and L. Stamford, "Net zero in the heating sector: Technological options and environmental sustainability from now to 2050," *Energy Conversion and Management*, vol. 230, p. 113838, Feb. 2021, doi: https://doi.org/10.1016/j.enconman.2021.113838.

[28] S. Mazumdar, D. Thakker, J. Hayes, N. Matos, and P. Bate, "Towards Achieving Net Zero by 2050 in the UK – Stakeholder Perspectives in Integrated Urban Planning," *Futures*, vol. 152, pp. 103197–103197, Sep. 2023, doi: https://doi.org/10.1016/j.futures.2023.103197.

[29] X. Li, H. Arbabi, G. Bennett, T. Oreszczyn, and D. Densley Tingley, "Net zero by 2050: Investigating carbon-budget compliant retrofit measures for the English housing stock,"

Renewable and Sustainable Energy Reviews, vol. 161, p. 112384, Jun. 2022, doi: https://doi.org/10.1016/j.rser.2022.112384.

[30] S. Fuger, A. Köhl, T. Stroehle, M. Lang, Johann Füller, and K. Hutter, "Fighting the wicked problem of plastic pollution and its consequences for developing regions with expert and crowd solutions," *Hawaii International Conference on System Sciences*, Jan. 2021, doi: https://doi.org/10.24251/hicss.2021.073.

[31] P. K. Mohanty, D. S. Roy, and K. H. K. Reddy, "Explainable AI for predicting daily household energy usages," *IEEE Xplore*, Dec. 01, 2022. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=10060217 (accessed Aug. 25, 2023).

[32] A. O. Ojo, Oladapo Tolulope Ibitoye, O. M. Adegoke, Tadese Victor A, K. C. J, and Oyeleye Ibrahim-Niass, "Design and Performance Analysis of Gradient-Based Optimization Techniques for Household Power Consumption Prediction," pp. 1–5, Nov. 2024, doi: https://doi.org/10.1109/nigercon62786.2024.10927052.

[33] H. Youssef, S. Kamel, M. H. Hassan, and L. Nasrat, "Optimizing energy consumption patterns of smart home using a developed elite evolutionary strategy artificial ecosystem optimization algorithm," *Energy*, vol. 278, p. 127793, Sep. 2023, doi: https://doi.org/10.1016/j.energy.2023.127793.

[34] Aakash Bhandary, Vruti Dobariya, Gokul Yenduri, R. H. Jhaveri, Saikat Gochhait, and F. Benedetto, "Enhancing Household Energy Consumption Predictions Through Explainable AI Frameworks," *IEEE Access*, vol. 12, pp. 36764–36777, Jan. 2024, doi: https://doi.org/10.1109/access.2024.3373552.

[35] Y. Ye, "Energy Consumption Management and Optimization of Artificial Intelligence Systems based on Computer Communication," *2025 3rd International Conference on*

Integrated Circuits and Communication Systems (ICICACS), pp. 1–6, Feb. 2025, doi: https://doi.org/10.1109/icicacs65178.2025.10968203.

[36] "Optimization strategy based on deep reinforcement learning for home energy management," *CSEE Journal of Power and Energy Systems*, 2020, doi: https://doi.org/10.17775/cseejpes.2019.02890.

[37] Y. He, J. Liu, and T. Zhang, "Home Energy Management System Optimization Strategy Based on Reinforcement Learning," pp. 24–30, Jan. 2021, doi: https://doi.org/10.1145/3448218.3448230.

[38] T.-Y. Kim and S.-B. Cho, "Particle Swarm Optimization-based CNN-LSTM Networks for Forecasting Energy Consumption," *2019 IEEE Congress on Evolutionary Computation (CEC)*, Jun. 2019, doi: https://doi.org/10.1109/cec.2019.8789968.

[39] T. M. Olatunde, A. C. Okwandu, D. O. Akande, and Z. Q. Sikhakhane, "REVIEWING THE ROLE OF ARTIFICIAL INTELLIGENCE IN ENERGY EFFICIENCY OPTIMIZATION," *Engineering Science & Technology Journal*, vol. 5, no. 4, pp. 1243–1256, Apr. 2024, doi: https://doi.org/10.51594/estj.v5i4.1015.

[40] M. Premkumar *et al.*, "Optimizing household energy management with distributed energy resources: A multi-learning-guided stochastic optimization approach," *Energy and Buildings*, vol. 331, p. 115323, Mar. 2025, doi: https://doi.org/10.1016/j.enbuild.2025.115323.

[41] Puja Gholap, "Comparative Analysis on Developed Optimization Techniques for Reducing Energy Consumption in AI Training and Inference," *Journal of Information Systems Engineering and Management*, vol. 10, no. 14s, pp. 111–116, Feb. 2025, doi: https://doi.org/10.52783/jisem.v10i14s.2166.

[42] K. Srikanth, “Unveiling Energy Insights: An Explainable AI-Driven Framework for Precision Household Consumption Forecasting,” *International Journal of Scientific Research and Engineering Trends*, vol. 11, no. 2, pp. 1400–1405, Mar. 2025, doi: https://doi.org/10.61137/ijsret.vol.11.issue2.231.

[43] J. Liu and Y. He, “Research on Energy Management and Optimization Algorithm Based on Artificial Intelligence in Energy-Saving Buildings,” pp. 985–990, May 2024, doi: https://doi.org/10.1109/telepe64216.2024.00183.

[44] N. Bassiliades and G. Chalkiadakis, “Artificial Intelligence Techniques for the Smart Grid,” *Advances in Building Energy Research*, vol. 12, no. 1, pp. 1–2, Dec. 2017, doi: https://doi.org/10.1080/17512549.2017.1410362.

[45] Q. Qiao, A. Yunusa-Kaltungo, and R. E. Edwards, “Developing a machine learning based building energy consumption prediction approach using limited data: Boruta feature selection and empirical mode decomposition,” *Energy Reports*, vol. 9, pp. 3643–3660, Dec. 2023, doi: https://doi.org/10.1016/j.egyr.2023.02.046.

[46] G. C. Nobre and E. Tavares, “Scientific literature analysis on big data and internet of things applications on circular economy: a bibliometric study,” *Scientometrics*, vol. 111, no. 1, pp. 463–492, Feb. 2017, doi: https://doi.org/10.1007/s11192-017-2281-6.

[47] Sunil Kumar Jauhar, S. Pratap, Lakshay, Sanjoy Kumar Paul, and Angappa Gunasekaran, “Internet of things based innovative solutions and emerging research clusters in circular economy,” *Operations Management Research*, vol. 16, no. 4, pp. 1968–1988, Oct. 2023, doi: https://doi.org/10.1007/s12063-023-00421-9.

[48] A. Cavalieri, J. Reis, and M. Amorim, “Circular Economy and Internet of Things: Mapping Science of Case Studies in Manufacturing Industry,” *Sustainability*, vol. 13, no. 6, p. 3299, Mar. 2021, doi: https://doi.org/10.3390/su13063299.

[49] Mohammad Nabil Almunawar, P. Ordóñez, and Muhammad Anshari, *Sustainable Development and the Digital Economy*. Taylor & Francis, 2023.

[50] Denisa-Roxana Botea-Muntean and R. Constantinescu, "Internet of Things – one of the Digital Steps that Companies Should Consider towards Circular Economy and Sustainability An Exploratory Case Study Defined by the Sustainable Retrospective of one of the Top Global Retail Groups," *Proceedings of the ... International Conference on Business Excellence*, vol. 17, no. 1, pp. 89–101, Jul. 2023, doi: https://doi.org/10.2478/picbe-2023-0011.

[51] R. Siegel, J. Antony, J. A. Garza-Reyes, A. Cherrafi, and B. Lameijer, "Integrated green lean approach and sustainability for SMEs: From literature review to a conceptual framework," *Journal of Cleaner Production*, vol. 240, p. 118205, Dec. 2019, doi: https://doi.org/10.1016/j.jclepro.2019.118205.

[52] C. P. C. Bong, Y. V. Fan, C. T. Lee, and J. J. Klemeš, "Process analysis and optimisation for a sustainable circular economy," *Cleaner Engineering and Technology*, vol. 11, p. 100578, Nov. 2022, doi: https://doi.org/10.1016/j.clet.2022.100578.

[53] K. Gopalakrishnan, "Deep Learning in Data-Driven Pavement Image Analysis and Automated Distress Detection: A Review," *Data*, vol. 3, no. 3, p. 28, Jul. 2018, doi: https://doi.org/10.3390/data3030028.

[54] A. Agrawal, K. Gopalakrishnan, and A. Choudhary, "Materials Image Informatics Using Deep Learning," *WORLD SCIENTIFIC eBooks*, pp. 205–230, Mar. 2020, doi: https://doi.org/10.1142/9789811204555_0006.

[55] R. Mohan, E. A. Gopalakrishnan, and K. P. Soman, "Classification of states of bi-stable oscillator using deep learning," *2017 International Conference on Advances in Computing,*

Communications and Informatics (ICACCI), pp. 1792–1796, Sep. 2017, doi: https://doi.org/10.1109/icacci.2017.8126104.

[56] C.-Y. Chen, J. Choi, K. Gopalakrishnan, V. Srinivasan, and S. Venkataramani, "Exploiting approximate computing for deep learning acceleration," *IEEE Xplore*, Mar. 01, 2018. https://ieeexplore.ieee.org/iel7/8337149/8341968/08342119.pdf?casa_token=RLMq-JSbJt0AAAAA:6LNVUIvasdGzT12qAv1yOO_KoeAj_A9DsapZWvdgjFJW8SUhxcoYQ6qwpjK97_Hu6ZeFcv076g (accessed Apr. 10, 2022).

[57] A. Gopalakrishnan, K. P. Soman, and B. Premjith, "A Deep Learning-Based Named Entity Recognition in Biomedical Domain," *Lecture notes in electrical engineering*, pp. 517–526, Jan. 2019, doi: https://doi.org/10.1007/978-981-13-5802-9_47.

[58] P. C. Roth, K. Huck, G. Gopalakrishnan, and F. Wolf, "Using Deep Learning for Automated Communication Pattern Characterization: Little Steps and Big Challenges," *Lecture notes in computer science*, pp. 265–272, Jan. 2019, doi: https://doi.org/10.1007/978-3-030-17872-7_16.

[59] R. Radha and R. Gopalakrishnan, "Use of Machine Learning and Deep Learning in Healthcare—A Review on Disease Prediction System," pp. 135–152, Jan. 2022, doi: https://doi.org/10.1002/9781119821908.ch6.

[60] D. Kohli, B. C. Das, V. Gopalakrishnan, and K. N. Iyer, "Learning rotation invariance in deep hierarchies using circular symmetric filters," pp. 2846–2850, Mar. 2017, doi: https://doi.org/10.1109/icassp.2017.7952676.

[61] P. Gopika, V. Sowmya, E. A. Gopalakrishnan, and K. P. Soman, "Performance Improvement of Residual Skip Convolutional Neural Network for Myocardial Disease Classification," *ICICCT 2019 – System Reliability, Quality Control, Safety, Maintenance and Management*, pp. 226–234, Jun. 2019, doi: https://doi.org/10.1007/978-981-13-8461-5_25.

[62] R. Gopalakrishnan, Y. Chua, and L. R. Iyer, “Classifying Neuromorphic Data Using a Deep Learning Framework for Image Classification,” *2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pp. 1520–1524, Nov. 2018, doi: https://doi.org/10.1109/icarcv.2018.8581256.

[63] G. K, “Machine Learning Models for Flood Severity Prediction,” *International Journal for Research in Applied Science and Engineering Technology*, vol. 9, no. 3, pp. 1424–1430, Mar. 2021, doi: https://doi.org/10.22214/ijraset.2021.33509.

[64] G. Zhang, L. Liu, and S. Song, “A review on the integration of renewable energy sources, energy storage and electric vehicles with smart grid,” *Sustainable Cities and Society*, vol. 65, p. 102573, 2021.

[65] J. Hu, Y. Wang, and H. Liu, “Dynamic energy management in microgrids using artificial intelligence,” *IEEE Access*, vol. 8, pp. 65384–65401, 2020.

[66] Ellen MacArthur Foundation, “Completing the picture: How the circular economy tackles climate change,” 2019.

[67] R. Geyer, J. R. Jambeck, and K. L. Law, “Production, use, and fate of all plastics ever made,” *Science Advances*, vol. 3, no. 7, p. e1700782, 2017.

[68] A. Sharma and J. Manhas, “A comprehensive review of artificial intelligence in smart grids,” *Energy Systems*, vol. 10, pp. 485–511, 2019.

[69] S. A. Pourmousavi and M. H. Nehrir, “Real-time central demand response for primary frequency regulation,” *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 1988–1995, 2012.

[70] M. S. Hossain, G. Muhammad, and M. Guizani, “Explainable AI and mass surveillance system-based healthcare framework to combat COVID-I9 like pandemics,” *IEEE Network*, vol. 34, no. 4, pp. 126–132, 2020.

[71] B. B. Zhou, H. Chen, and Y. Xie, “AI-based predictive models for energy management

in buildings: A review," *Renewable and Sustainable Energy Reviews*, vol. 149, p. 111333, 2021.

[72] A. Y. Saber and G. K. Venayagamoorthy, "Plug-in vehicles and renewable energy sources for cost and emission reductions," *IEEE Trans. Ind. Electron.*, vol. 58, no. 4, pp. 1229–1238, 2011.

[73] M. Hannan, A. Lipu, A. Hussain, and A. Mohamed, "A review of energy management system for electric vehicles," *Renewable and Sustainable Energy Reviews*, vol. 78, pp. 1389–1405, 2017.

[74] A. Sharma et al., "Waste collection route optimisation using AI-based smart bins," *IEEE Trans. Ind. Inform.*, vol. 17, no. 4, pp. 2738–2747, 2021.

[75] M. A. Ponce-Cueto, R. A. Saade, and C. Gomez, "Modelling of waste-to-energy plants using machine learning," *Waste Management*, vol. 120, pp. 270–278, 2021.

[76] L. Song, J. Zhou, and H. Wang, "Intelligent robotic sorting system for solid waste recycling," *IEEE Access*, vol. 9, pp. 32412–32421, 2021.

[77] T. Silva, R. You, and R. González, "An overview of AI technologies applied to circular economy," *Procedia CIRP*, vol. 90, pp. 730–735, 2020.

[78] S. Ahmed, B. Yildiz, and R. Merlet, "Challenges in explainable AI for energy systems," *Energy and AI*, vol. 3, p. 100041, 2021.

[79] Y. Wang et al., "Hybrid AI model for power load prediction in smart grids," *IEEE Trans. Ind. Appl.*, vol. 56, no. 2, pp. 1478–1485, 2020.

[80] F. J. Díaz-González, A. Sumper, and R. Villafáfila-Robles, "Review of predictive models for integrated energy systems," *Applied Energy*, vol. 135, pp. 232–244, 2014.

[81] L. Mora, J. R. Deakin, and N. Reid, "Smart city development paths: Insights from 22 global strategies," *Cities*, vol. 92, pp. 45–61, 2019.

[82] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller, "Causability and

explainability of AI in medicine," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 9, no. 4, p. e1312, 2019.

[83] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable AI: A review of machine learning interpretability methods," *Entropy*, vol. 23, no. 1, p. 18, 2021.

[84] T. Ahmad, M. Chen, and A. Wang, "A review on machine learning forecasting techniques for energy load prediction," *Renewable and Sustainable Energy Reviews*, vol. 134, p. 110255, 2020.

[85] H. Liu, Y. Tian, and K. Yu, "A deep learning approach for building energy consumption prediction," *Applied Energy*, vol. 240, pp. 263–274, 2019.

[86] G. Dudek, "Short-term load forecasting using random forests," *Electric Power Systems Research*, vol. 81, no. 3, pp. 1444–1451, 2011.

[87] Y. Chen, W. Wang, and R. Huang, "Load forecasting using support vector machines with feature selection," *IEEE Trans. Power Syst.*, vol. 23, no. 4, pp. 1688–1698, 2008.

[88] T. Hong, P. Wang, and H. Luo, "State-of-the-art short-term load forecasting techniques," *Electric Power Components and Systems*, vol. 42, no. 10, pp. 1033–1044, 2014.

[89] A. Marino, D. Santamaria, and F. Benzi, "LSTM networks for forecasting of energy demand and price in smart grids," *Energies*, vol. 12, no. 13, p. 2533, 2019.

[90] L. Chowdhury, M. Ferdous, and R. Islam, "Prediction of municipal solid waste generation using ML algorithms: A case study of Bangladesh," *Waste Management*, vol. 118, pp. 593–605, 2020.

[91] C. Silva, H. Cunha, and J. Brito, "GIS-based decision support system for solid waste collection optimization," *Waste Management*, vol. 34, no. 4, pp. 875–886, 2014.

[92] F. Ferronato and V. Torretta, "Waste mismanagement in developing countries: A review of global issues," *International Journal of Environmental Research and Public Health*, vol. 16, no. 6, p. 1060, 2019.

[93] A. Sharma and J. Kumar, “Smart bin architecture and AI-based predictive waste collection,” *IEEE Access*, vol. 8, pp. 105507–105521, 2020.

[94] N. Kumar and P. Jain, “Forecasting sustainability KPIs with AI: A review of global practices,” *Sustainability,* vol. 13, no. 2, p. 483, 2021.

[95] E. Giarola, C. G. Lemus-Zúñiga, and A. Hawkes, “Sustainable energy transitions: Using AI to evaluate policy trade-offs,” *Renewable Energy,* vol. 175, pp. 688–701, 2021.

[96] P. J. Antsaklis, “On intelligent control,” *IEEE Control Systems Magazine,* vol. 11, no. 3, pp. 5–13, 1991.

[97] M. Suganthi and A. A. Samuel, “Energy models for demand forecasting—A review,” *Renewable and Sustainable Energy Reviews,* vol. 16, no. 2, pp. 1223–1240, 2012.

[98] A. Holzinger et al., “Causability and explainability of artificial intelligence in medicine,” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery,* vol. 9, no. 4, p. e1312, 2019.

[99] A. Tjoa and C. Guan, “A survey on explainable artificial intelligence (XAI): Towards medical XAI,” *IEEE Trans. Neural Networks Learn. Syst.*, vol. 32, no. 11, pp. 4793–4813, 2021.

[100] M. Ribeiro, S. Singh, and C. Guestrin, “Why should I trust you?”: Explaining the predictions of any classifier,” in *Proc. 22nd ACM SIGKDD,* 2016, pp. 1135–1144.

[101] S. Lundberg and S. Lee, “A unified approach to interpreting model predictions,” in *Proc. NIPS*, 2017, pp. 4765–4774.

[102] F. Florez and K. Singh, “A survey of hybrid AI techniques in sustainable infrastructure systems,” *Journal of Environmental Informatics,* vol. 37, no. 3, pp. 154–167, 2021.

[103] K. Mahapatra, B. Banerjee, and R. K. Roy, “Integration of AI and lifecycle assessment in urban sustainability,” *Sustainable Cities and Society,* vol. 75, p. 103288, 2021.

[104] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Pearson, 2010.

[105] M. Negnevitsky, *Artificial Intelligence: A Guide to Intelligent Systems*, 3rd ed., Addison-Wesley, 2011.

[106] Y. Wang, Q. Chen, C. Kang, and Q. Xia, “Incorporating operational constraints and user behavior into electric vehicle energy management,” *IEEE Trans. Smart Grid*, vol. 5, no. 3, pp. 1421–1429, 2014.

[107] A. W. Colombo et al., “Industrial automation based on cyber-physical systems technologies: Prototype implementations and challenges,” *Comput. Ind.*, vol. 81, pp. 11–25, 2016.

[108] A. Elkamel, A. Albahri, and K. Fatemi, “Optimization of waste collection using GIS,” *Journal of Environmental Informatics*, vol. 21, no. 1, pp. 18–28, 2013.

[109] J. Johansson and M. Hellström, “The effect of dynamic routing in urban waste collection systems,” *Waste Management*, vol. 33, no. 4, pp. 833–839, 2013.

[110] G. Dudek, “Short-term load forecasting using random forests,” *Electric Power Systems Research*, vol. 81, no. 3, pp. 1444–1451, 2011.

[111] T. Ahmad, M. Chen, and A. Wang, “A review on machine learning forecasting techniques for energy load prediction,” *Renewable and Sustainable Energy Reviews*, vol. 134, p. 110255, 2020.

[112] H. Liu, Y. Tian, and K. Yu, “A deep learning approach for building energy consumption prediction,” *Applied Energy*, vol. 240, pp. 263–274, 2019.

[113] B. Dong, Z. Cao, and S. E. Lee, “Applying support vector machines to predict building energy consumption in tropical region,” *Energy and Buildings*, vol. 37, no. 5, pp. 545–553, 2005.

[114] M. Marino et al., “LSTM-based deep learning models for short-term electric load

forecasting," *Energies*, vol. 12, no. 1, p. 147, 2019.

[115] J. Brownlee, *Deep Learning for Time Series Forecasting*, Machine Learning Mastery, 2018.

[116] M. Hannan, A. Lipu, A. Hussain, and A. Mohamed, "A review of energy management system for electric vehicles," *Renewable and Sustainable Energy Reviews*, vol. 78, pp. 1389–1405, 2017.

[117] A. Sharma and J. Kumar, "Smart bin architecture and AI-based predictive waste collection," *IEEE Access*, vol. 8, pp. 105507–105521, 2020.

[118] C. Silva, H. Cunha, and J. Brito, "GIS-based decision support system for solid waste collection optimization," *Waste Management*, vol. 34, no. 4, pp. 875–886, 2014.

[119] F. Florez and K. Singh, "A survey of hybrid AI techniques in sustainable infrastructure systems," *Journal of Environmental Informatics*, vol. 37, no. 3, pp. 154–167, 2021.

[120] M. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?": Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD*, 2016, pp. 1135–1144.

[121] A. Holzinger et al., "Causability and explainability of artificial intelligence in medicine," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 9, no. 4, p. e1312, 2019.

[122] A. Tjoa and C. Guan, "A survey on explainable artificial intelligence (XAI): Towards medical XAI," *IEEE Trans. Neural Networks Learn. Syst.*, vol. 32, no. 11, pp. 4793–4813, 2021.

[123] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable AI: A review of machine learning interpretability methods," *Entropy*, vol. 23, no. 1, p. 18, 2021.

[124] K. Mahapatra, B. Banerjee, and R. K. Roy, "Integration of AI and lifecycle assessment in urban sustainability," *Sustainable Cities and Society*, vol. 75, p. 103288, 2021.

[125] E. Giarola et al., "Sustainable energy transitions: Using AI to evaluate policy trade-

offs," *Renewable Energy*, vol. 175, pp. 688–701, 2021.

[126] F. Florez et al., "Data-efficient machine learning for sustainable systems," *Journal of Cleaner Production*, vol. 256, p. 120366, 2020.

[127] N. Kumar and P. Jain, "Forecasting sustainability KPIs with AI: A review of global practices," *Sustainability*, vol. 13, no. 2, p. 483, 2021.

[128] D. Santamaria and A. Marino, "Benchmarking rule-based vs AI models in energy systems," *Energy Reports*, vol. 7, pp. 884–899, 2021.

[129] A. Graves, "Supervised sequence labelling with recurrent neural networks," *Studies in Computational Intelligence*, vol. 385, 2012.

[130] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.

[131] S. Hippert, C. Pedreira, and R. Souza, "Neural networks for short-term load forecasting: A review and evaluation," *IEEE Trans. Power Syst.*, vol. 16, no. 1, pp. 44–55, 2001.

[132] T. Ahmad et al., "A review on machine learning forecasting techniques for energy load prediction," *Renew. Sustain. Energy Rev.*, vol. 134, p. 110255, 2020.

[133] J. Kong et al., "Short-term residential load forecasting based on LSTM recurrent neural network," *IEEE Trans. Smart Grid*, vol. 10, no. 1, pp. 841–851, 2019.

[134] H. Liu and Y. Tian, "A deep learning approach for building energy consumption prediction," *Applied Energy*, vol. 240, pp. 263–274, 2019.

[135] A. Marino, D. Santamaria, and F. Benzi, "LSTM networks for forecasting of energy demand and price in smart grids," *Energies*, vol. 12, no. 13, p. 2533, 2019.

[136] W. Wang et al., "Hybrid CNN–LSTM model for short-term load forecasting," *Energy Reports*, vol. 7, pp. 115–125, 2021.

[137] X. Zhang, J. Xu, and B. Wang, "Attention-based LSTM for solar power forecasting," *Renewable Energy*, vol. 180, pp. 383–393, 2021.

[138] T. Hong and S. Fan, “Probabilistic electric load forecasting: A tutorial review,” *Int. J. Forecast.*, vol. 32, no. 3, pp. 914–938, 2016.

[139] J. Qdais, H. Hamoda, and M. Newham, “Analysis of residential solid waste at generation sites,” *Waste Management and Research*, vol. 15, no. 4, pp. 395–406, 1997.

[140] R. Geyer, J. Jambeck, and K. Law, “Production, use, and fate of all plastics ever made,” *Science Advances*, vol. 3, no. 7, p. e1700782, 2017.

[141] L. Chowdhury, M. Ferdous, and R. Islam, “Prediction of municipal solid waste generation using ML algorithms,” *Waste Management*, vol. 118, pp. 593–605, 2020.

[142] A. Sharma and J. Kumar, “Smart bin architecture and AI-based predictive waste collection,” *IEEE Access*, vol. 8, pp. 105507–105521, 2020.

[143] M. S. Hossain et al., “Urban waste monitoring using deep learning and IoT,” *Sensors*, vol. 19, no. 24, p. 5585, 2019.

[144] Y. Liu and L. Zheng, “Deep learning prediction of landfill gas emissions using environmental time-series data,” *Journal of Environmental Management*, vol. 277, p. 111426, 2021.

[145] R. Islam and S. Hoque, “Waste-to-energy systems in developing countries: A case for AI-enhanced process modelling,” *Energy*, vol. 200, p. 117453, 2020.

[146] A. Holzinger et al., “Causability and explainability of AI in medicine,” *Wiley Interdiscip. Rev. Data Mining Knowl. Discov.*, vol. 9, no. 4, p. e1312, 2019.

[147] M. Ribeiro, S. Singh, and C. Guestrin, “Why should I trust you?”: Explaining the predictions of any classifier,” in *Proc. ACM SIGKDD*, 2016, pp. 1135–1144.

[148] S. Ahmed, B. Yildiz, and R. Merlet, “Challenges in explainable AI for energy systems,” *Energy and AI*, vol. 3, p. 100041, 2021.

[149] K. Mahapatra and R. K. Roy, “Hybrid modelling for urban energy and waste flows,” *Sustainable Cities and Society*, vol. 85, p. 104049, 2022.

[150] F. Doshi-Velez and B. Kim, “Towards a rigorous science of interpretable machine learning,” arXiv:1702.08608, 2017.

[151] A. Adadi and M. Berrada, “Peeking inside the black-box: A survey on explainable artificial intelligence (XAI),” *IEEE Access*, vol. 6, pp. 52138–52160, 2018.

[152] S. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proc. NIPS*, 2017, pp. 4765–4774.

[153] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should I trust you?”: Explaining the predictions of any classifier,” in *Proc. ACM SIGKDD*, 2016, pp. 1135–1144.

[154] U. Mahajan et al., “Preserving fairness in predictive models using explainable AI,” *arXiv preprint arXiv:1910.06177*, 2019.

[155] K. Holzinger et al., “Causability and explainability of AI in medicine,” *Wiley Interdiscip. Rev. Data Mining Knowl. Discov.*, vol. 9, no. 4, p. e1312, 2019.

[156] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, “Explainable AI: A review of machine learning interpretability methods,” *Entropy*, vol. 23, no. 1, p. 18, 2021.

[157] T. Ribeiro, A. Prates, and B. Barbosa, “Explainable energy predictions for photovoltaic systems using LIME and SHAP,” *Applied Energy*, vol. 290, p. 116726, 2021.

[158] L. Chowdhury, M. Ferdous, and R. Islam, “Prediction of municipal solid waste generation using ML algorithms,” *Waste Management*, vol. 118, pp. 593–605, 2020.

[159] C. Silva, H. Cunha, and J. Brito, “GIS-based decision support system for waste collection optimisation,” *Waste Management*, vol. 34, no. 4, pp. 875–886, 2014.

[160] M. Tjoa and C. Guan, “A survey on explainable artificial intelligence (XAI): Towards medical XAI,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 11, pp. 4793–4813, 2021.

[161] D. Gunning et al., “XAI—Explainable Artificial Intelligence,” *DARPA*, 2018.

[162] K. Mahapatra, B. Banerjee, and R. Roy, “Integration of AI and lifecycle assessment in urban sustainability,” *Sustainable Cities and Society*, vol. 75, p. 103288, 2021.

[163] F. Florez and K. Singh, "A survey of hybrid AI techniques in sustainable infrastructure systems," *J. Environ. Informatics*, vol. 37, no. 3, pp. 154–167, 2021.