Agentic AI Energy Management: LLM-Enhanced Decision-Making in Battery Energy Systems

Dr Zekun Guo
Lecturer
Data Science, Artificial Intelligence, and Modelling (DAIM) Centre
University of Hull

Transformers and Self-Attention

"Attention Is All You Need"

- Q (Query), K (Key), V (Value)
- Generative pre-trained transformer (GPT)
- AI inference and reasoning

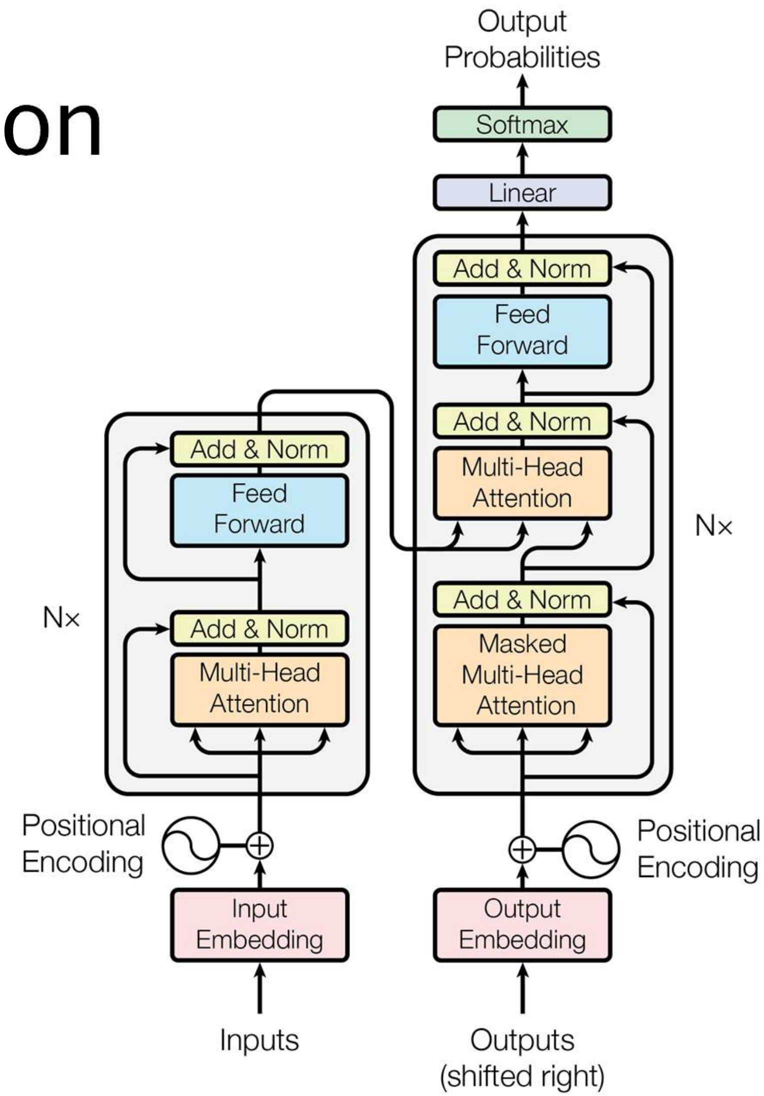

Vaswani et al (2018) **Attention is all you need**. Proc. of NeurIPS. The Transformer- model architecture [1]

ReAct Agents

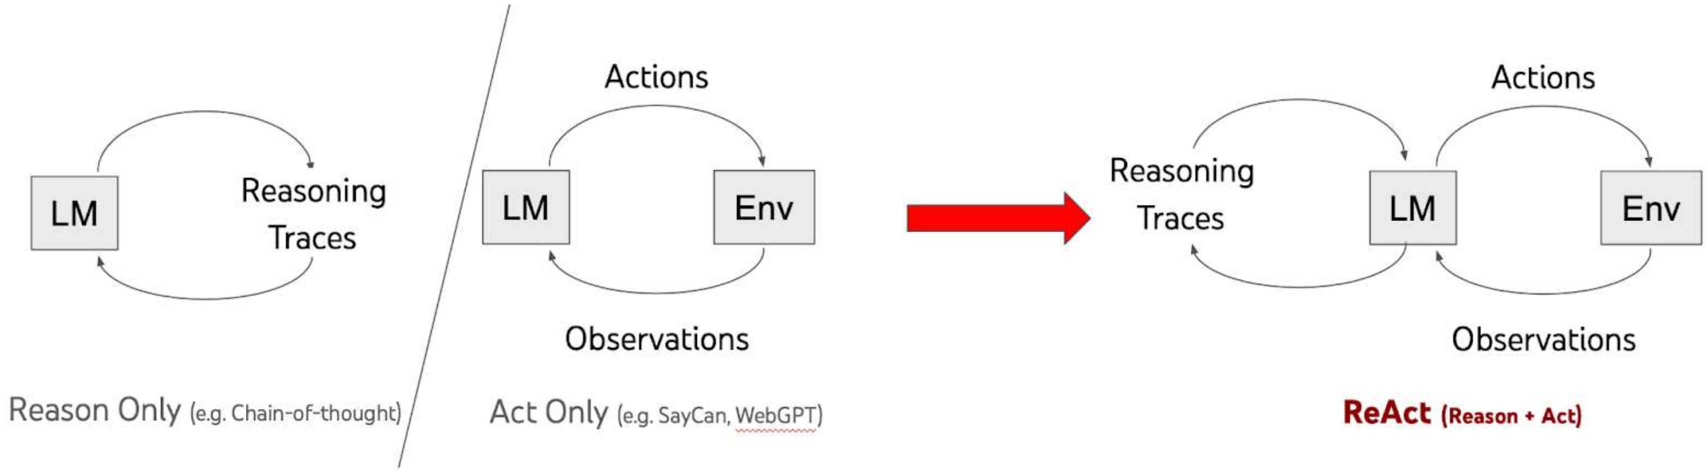

Combining large language models' reasoning (chain-of-thought prompting) and acting capabilities [2]

Yao, S et al (2023). **React: Synergizing reasoning and acting in language *models*.** *ICLR 2023.*

Multi-agent Frameworks – Camel AI

Camel AI

Other Multi-agent Frameworks

MetaGPT LangGraph CrewAI

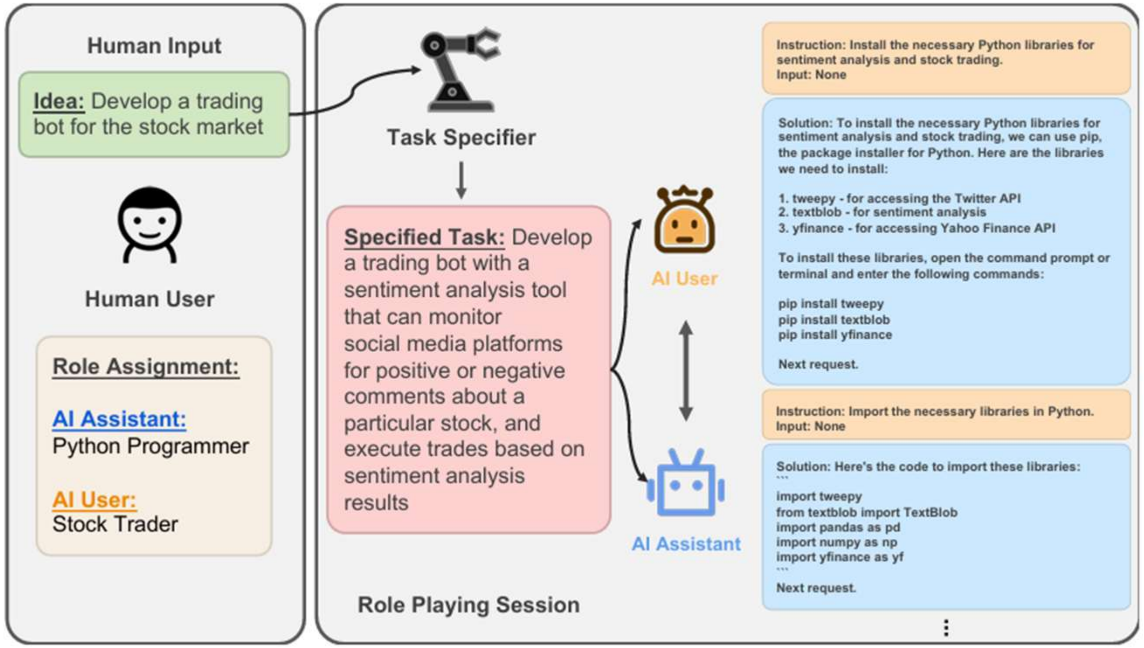

CAMEL Role-Playing Framework. [3]

Li, G et al (2023). **CAMEL: Communicative Agents for "Mind" Exploration of Large Language Model Society**. *NeurIPS 2023.*

Agentic Battery Energy Storage Management

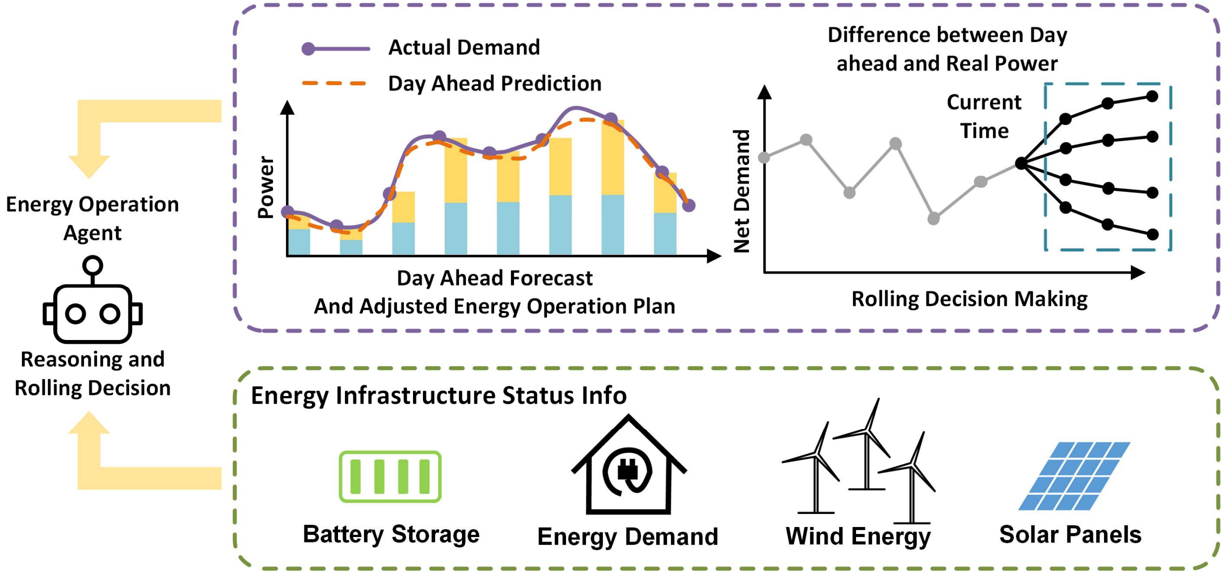

Diagram of Agentic Energy Operation Management

- **System Message Role:** "You are a power system dispatch optimization expert" overseeing real-time adjustments.
- **Day-Ahead Plan:** Generate 24-hour forecasts for load, wind, and PV.
- **Rolling Optimization:** Execute a real-time rolling horizon optimization using fixed window intervals
- **Parameter Adjustment:** Dynamically update settings like objective function and/or constraints.

Case Study

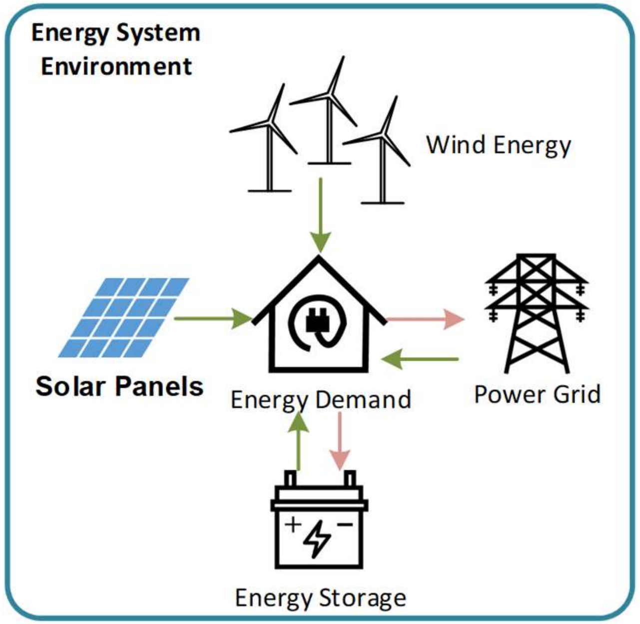

- Battery Energy Storage: 400 kWh (60 kW charge/discharge), initialized at 200 kWh, operating between 20% and 85% SOC.

- Integrated renewable generation from wind and photovoltaic (PV) systems, 250 kW, respectively.

- Grid connection featuring specific electricity purchase and sale prices that influence the dispatch cost.

- A day-ahead optimisation generates a 24-hour plan and then there is a 96 fifteen-minute intervals rolling intraday optimisation, serving as the reference for real-time operations.

Case Study

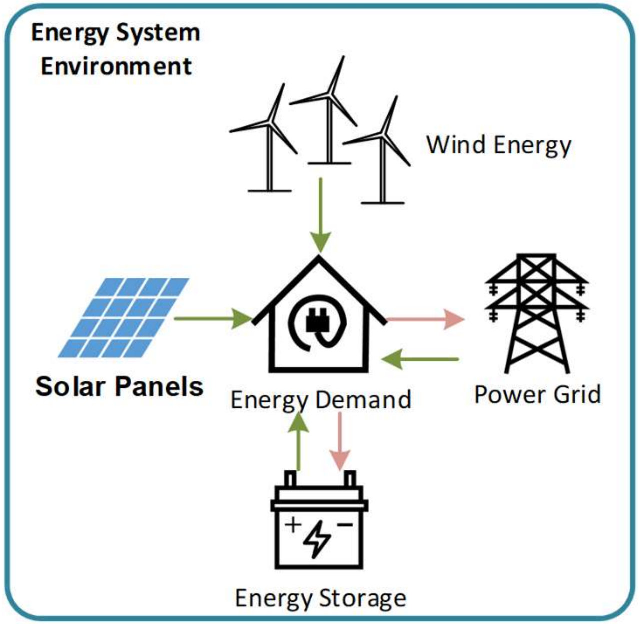

Energy System
Environment
Wind Energy
Solar Panels
Energy Demand
Power Grid
Energy Storage

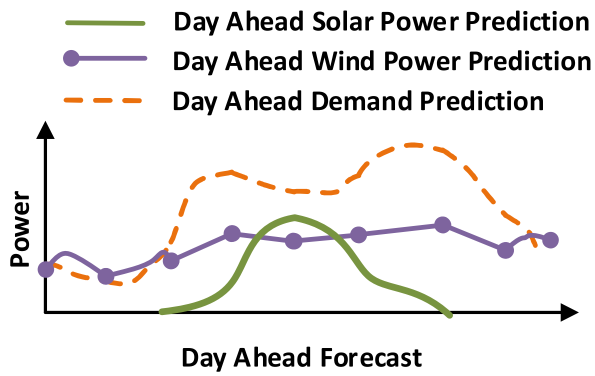

Day Ahead Solar Power Prediction
Day Ahead Wind Power Prediction
Day Ahead Demand Prediction
Power
Day Ahead Forecast

Case Study

Step 1:
Day-ahead
Optimisation

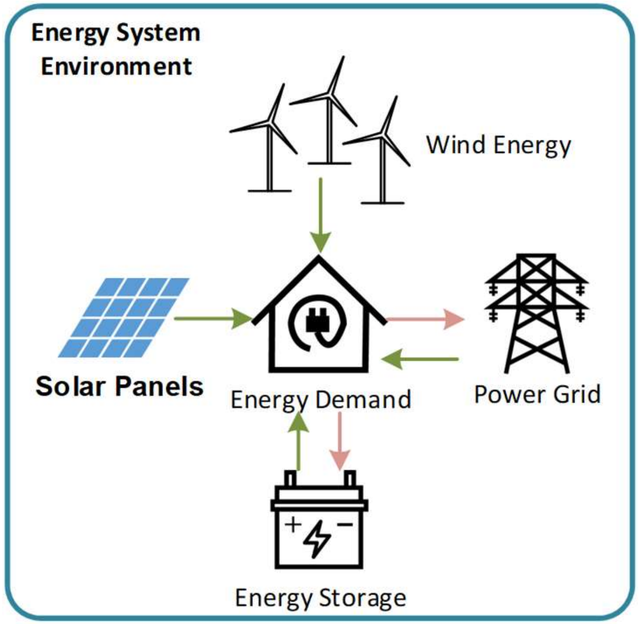

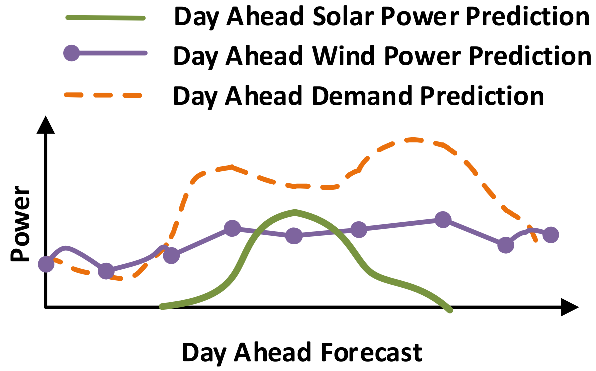

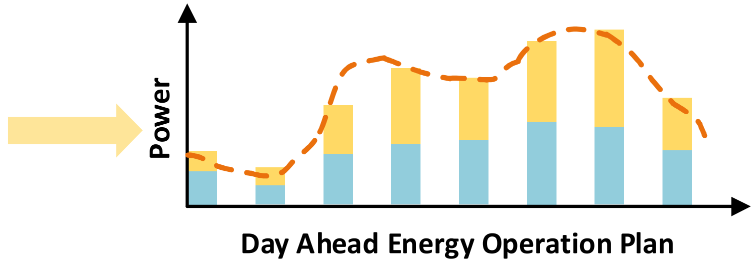

Case Study

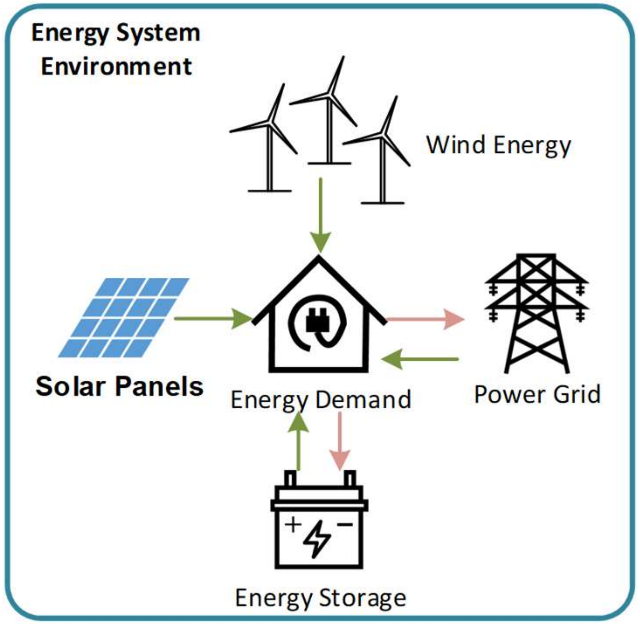

Energy System Environment
Wind Energy
Solar Panels
Energy Demand
Power Grid
Energy Storage

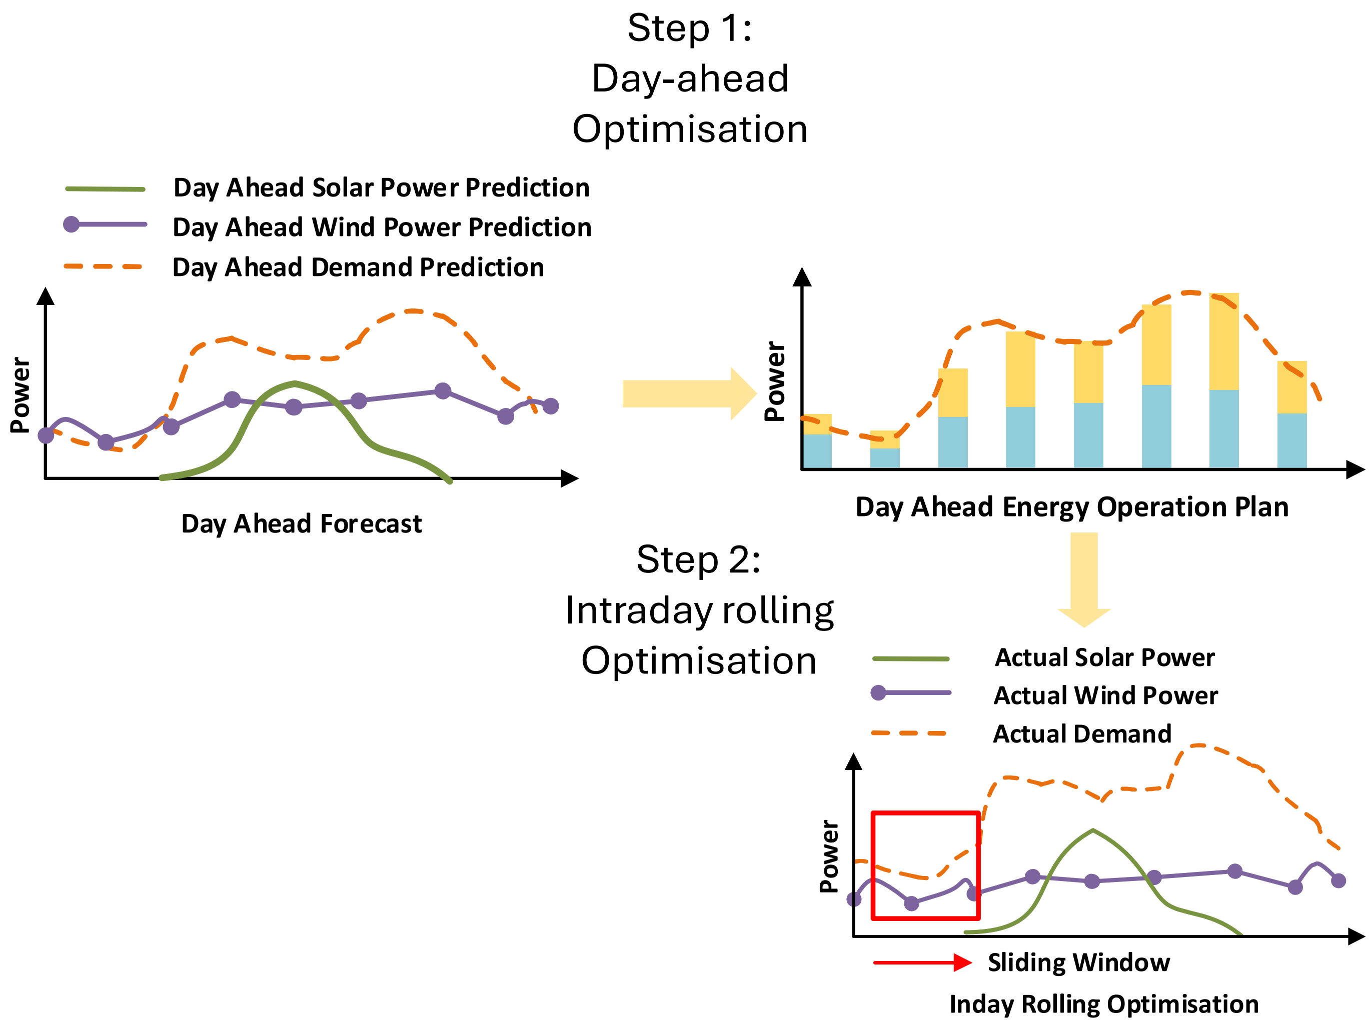

Step 1:
Day-ahead
Optimisation
Day Ahead Solar Power Prediction
Day Ahead Wind Power Prediction
Day Ahead Demand Prediction
Power
Day Ahead Forecast
Power
Day Ahead Energy Operation Plan
Step 2:
Intraday rolling
Optimisation
Actual Solar Power
Actual Wind Power
Actual Demand
Power
Sliding Window
Inday Rolling Optimisation

Case Study

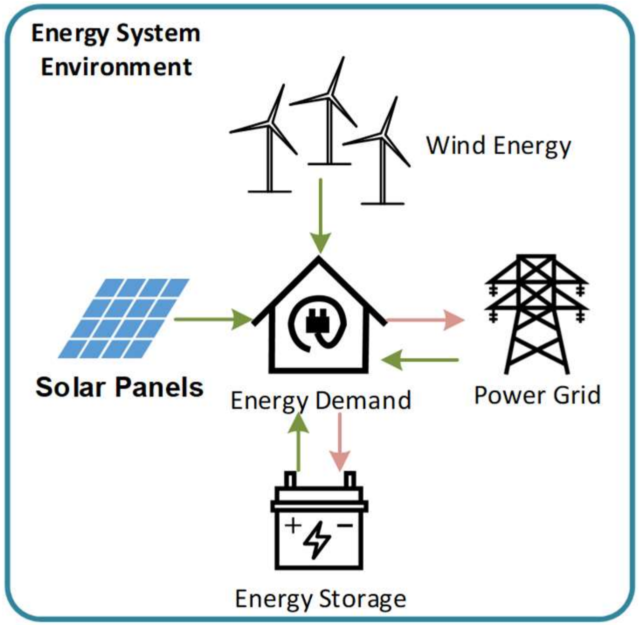

Energy System Environment
Wind Energy
Solar Panels
Energy Demand
Power Grid
Energy Storage

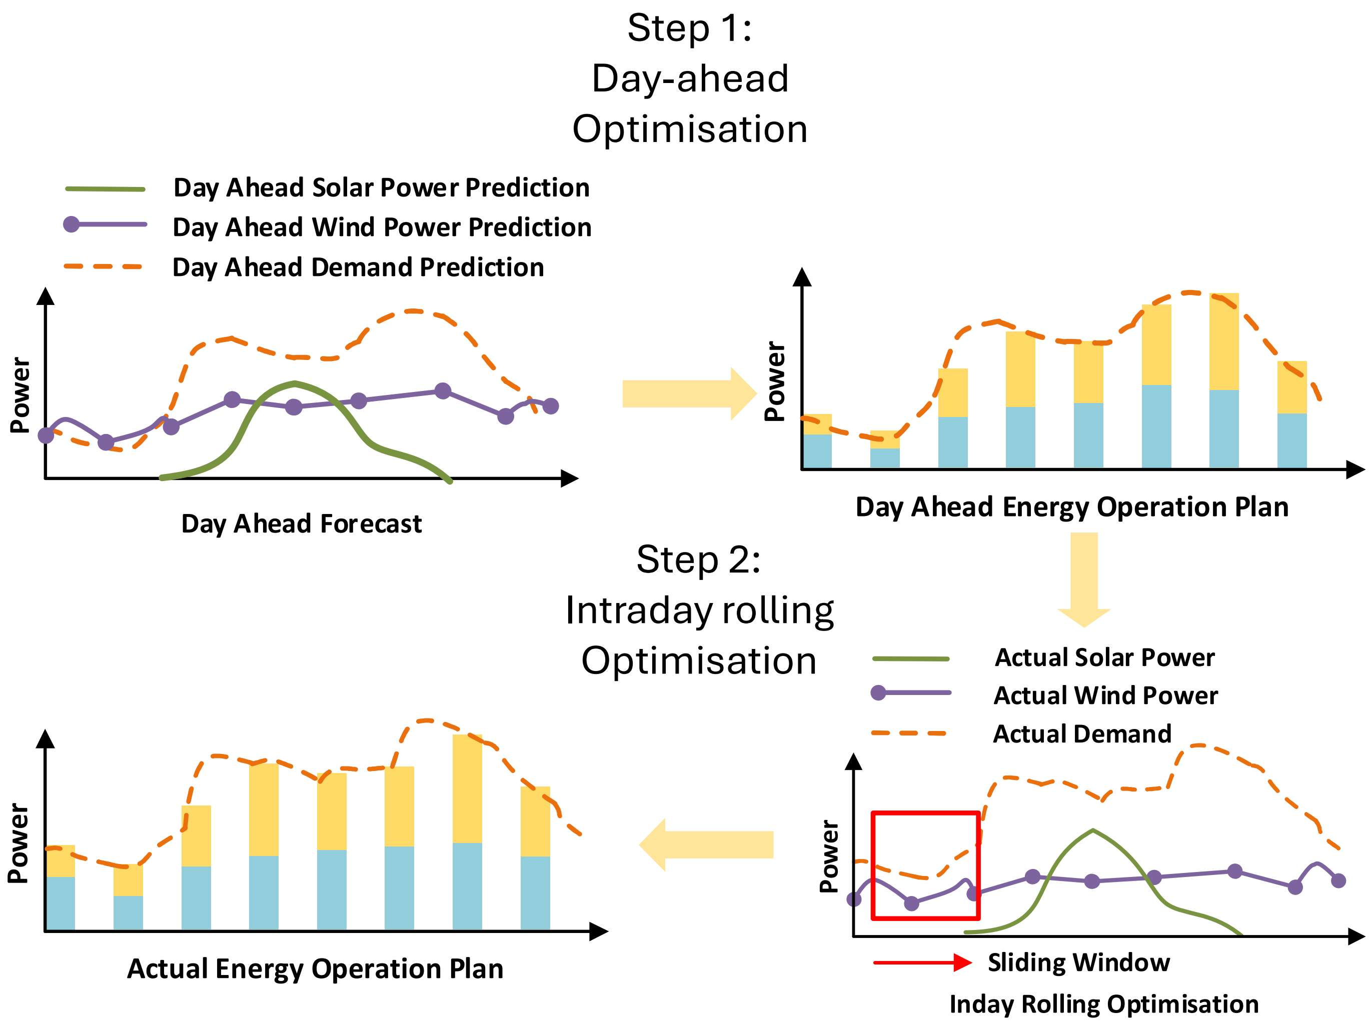

Step 1:
Day-ahead
Optimisation
Day Ahead Solar Power Prediction
Day Ahead Wind Power Prediction
Day Ahead Demand Prediction
Power
Day Ahead Forecast
Power
Day Ahead Energy Operation Plan
Step 2:
Intraday rolling
Optimisation
Actual Solar Power
Actual Wind Power
Actual Demand
Power
Sliding Window
Inday Rolling Optimisation
Power
Actual Energy Operation Plan

Case Study

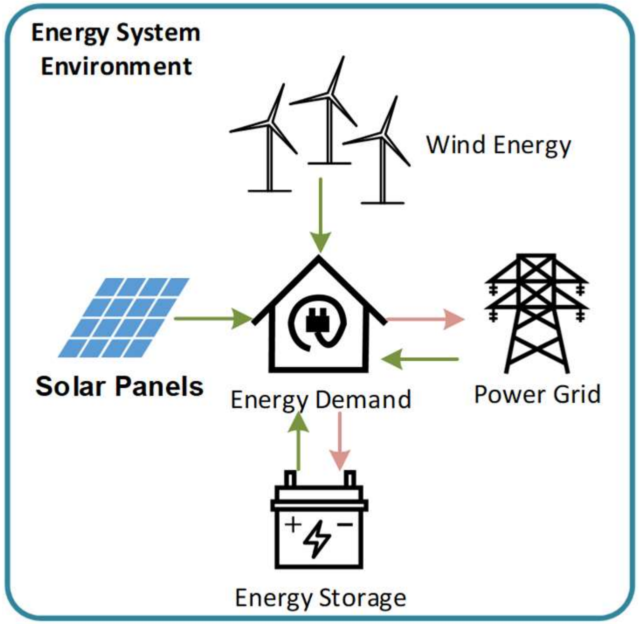
Energy System Environment
Wind Energy
Solar Panels
Energy Demand
Power Grid
Energy Storage

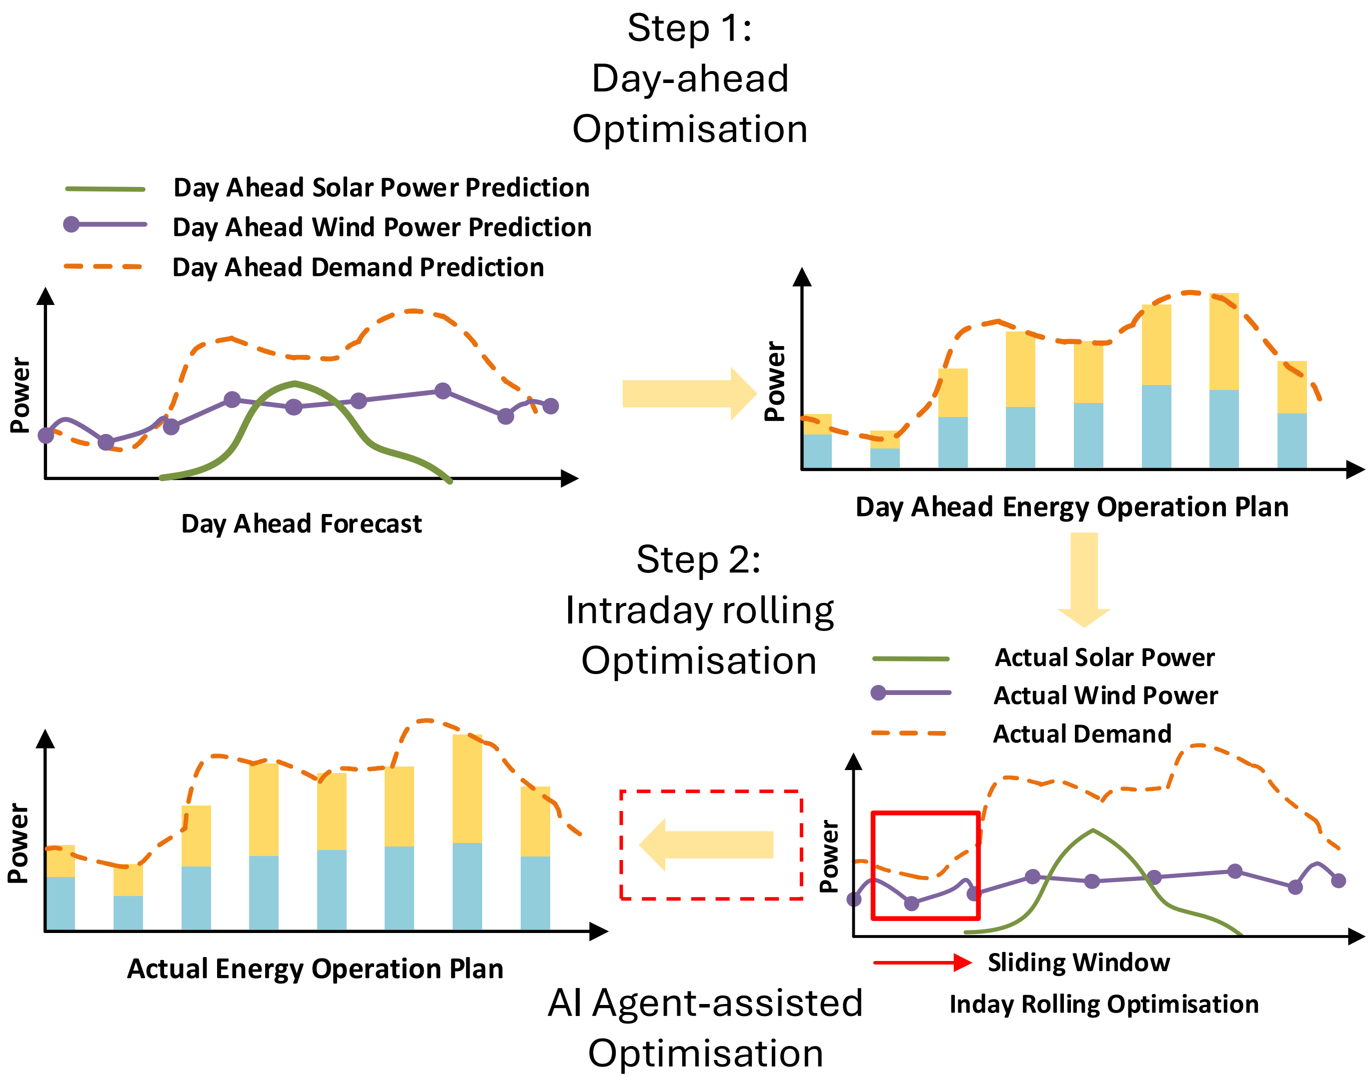
Step 1:
Day-ahead
Optimisation
Day Ahead Solar Power Prediction
Day Ahead Wind Power Prediction
Day Ahead Demand Prediction
Power
Day Ahead Forecast
Power
Day Ahead Energy Operation Plan
Step 2:
Intraday rolling
Optimisation
Actual Solar Power
Actual Wind Power
Actual Demand
Power
Sliding Window
Inday Rolling Optimisation
Power
Actual Energy Operation Plan
AI Agent-assisted
Optimisation

Prompts and Input Information

Test 0 – Fixed Penalty Coefficient Intraday Rolling Optimization

Baseline intraday rolling optimization that uses a fixed penalty coefficient (e.g., 0.01). The day-ahead 24-hour plan is expanded to 96 intervals, and a moving 4-hour window is used to minimise deviations between actual operations and the reference schedule.

Agents for all tests are based on OpenAI GPT-4o.

System Message:
You are a power system dispatch optimization expert.
The system consists of:
- Battery storage (400 kWh; initial 200 kWh; 20%-85% SOC; 60 kW)
- Wind and PV generation forecasts
- Grid connection with given buy/sell prices
- Day-ahead plan (24h) expanded to 96 pts

Task: Adjust the penalty coefficient based on current metrics. (Response in JSON)

Rolling Prompts:
Current Window Metrics:
- Start Time index t
- Current penalty (e.g., 0.01)

Upcoming Forecast:
- Ref. net grid exchange for next 16 intervals
- Forecasted load
- Forecasted renewable generation

Historical Data:
- Past actual vs. Ref.
- Historical penalty
- Average deviation

Test 1 - Dynamic Penalty Coefficient Adjustment Using LLM Assistance

System Message:
You are a power system dispatch optimization expert.
The system consists of:
- Battery storage (400 kWh; initial 200 kWh; 20%-85% SOC; 60 kW)
- Wind and PV generation forecasts
- Grid connection with given buy/sell prices
- Day-ahead plan (24h) expanded to 96 pts

Task: Based on current performance metrics, recommend terminal battery SOC constraint for the next rolling window.(Response in JSON)

Rolling Prompts:
Current Window Metrics:
- Start Time index t
- Current battery SOC: e.g., 200 kWh
- Day ahead terminal SOC target for t

Upcoming Forecast:
- Ref. net grid exchange for next 16 intervals
- Forecasted load
- Forecasted renewable generation

Historical Data:
- Past actual vs. Ref.
- Historical penalty
- Average deviation

Test 2 - Terminal State-of-Charge (SOC) Constraint Adjustment with LLM Guidance

Performance

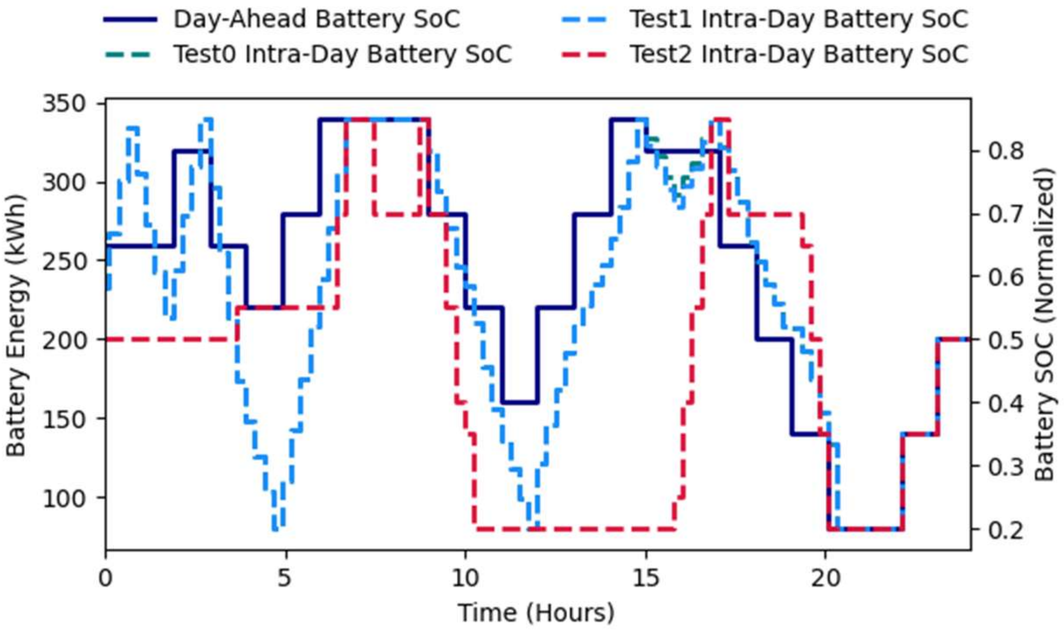

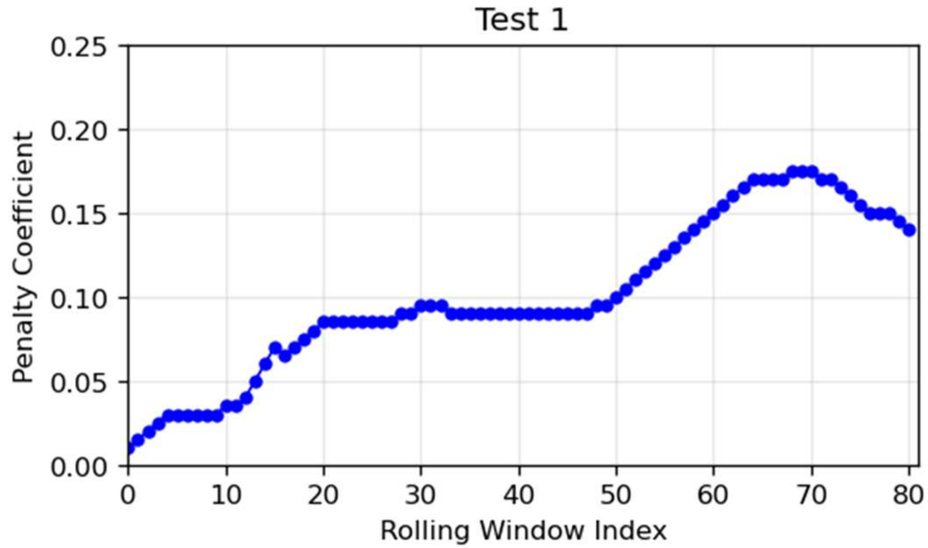

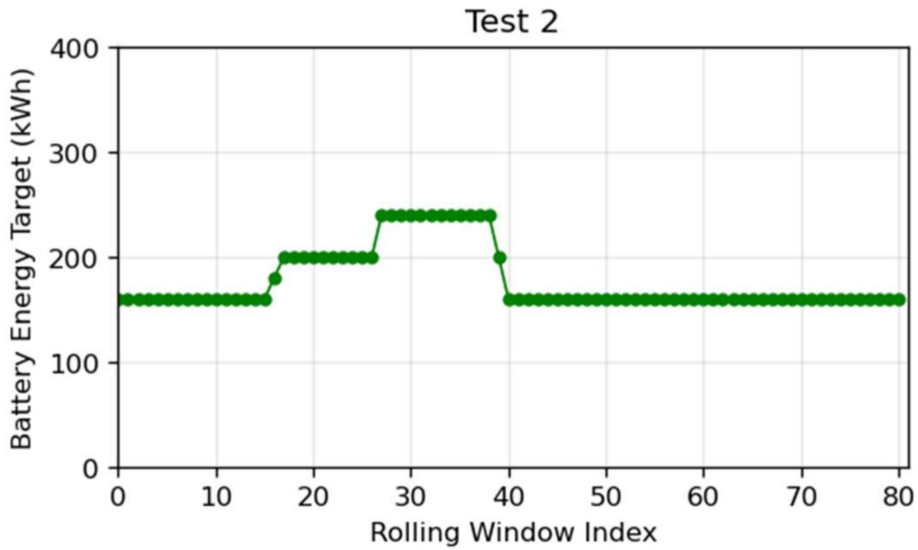

Inferences – Test 1 Dynamic Penalty Adjustment

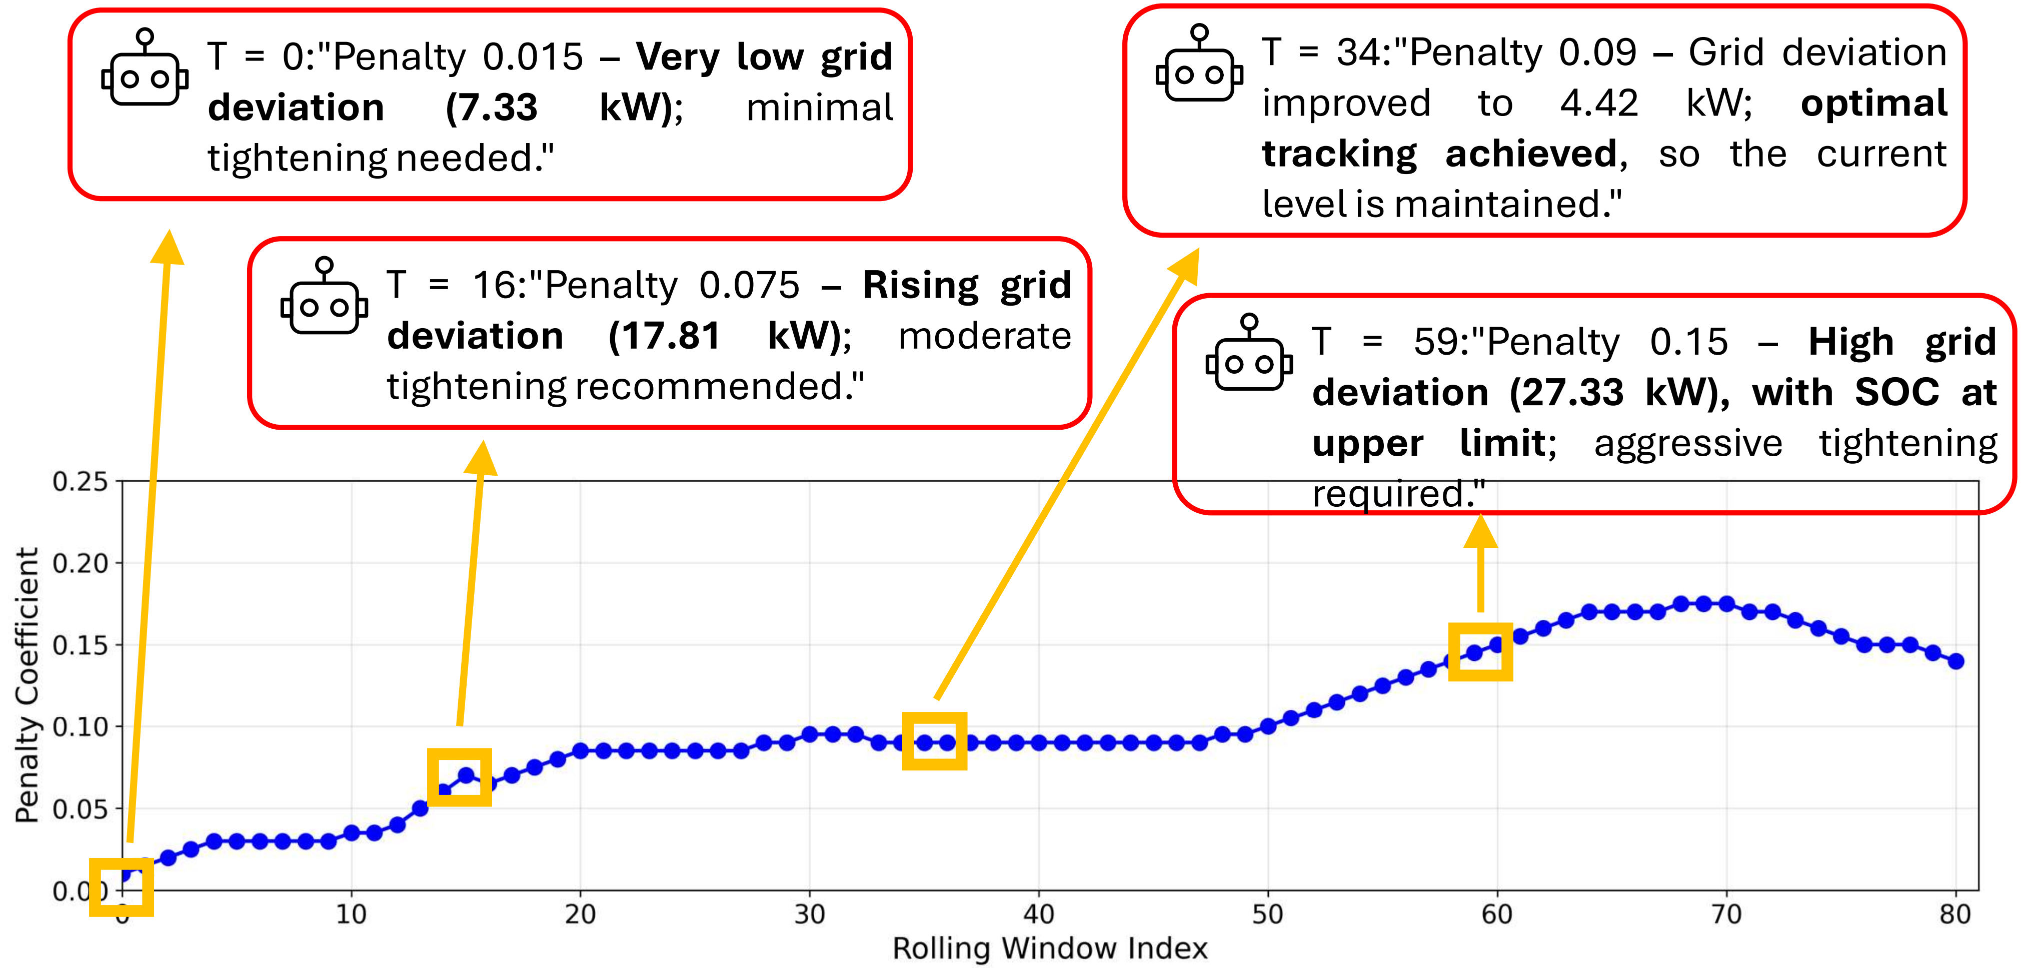

Inferences – Test 2 SOC Constraint Adjustment

T = 0:"Initial window – Terminal SOC remains 160 kWh; **stable forecast**, reserve target met."

T = 16:"Early increase – Terminal SOC raised to 200 kWh due to **high forecasted load**."

T = 26:"Mid-window surge – Terminal SOC bumped to 240 kWh; **strong load demand** requiring extra reserves."

T = 40:"Transition phase – Terminal Battery Energy drops back to 160 kWh to **restore minimum reserve level**."

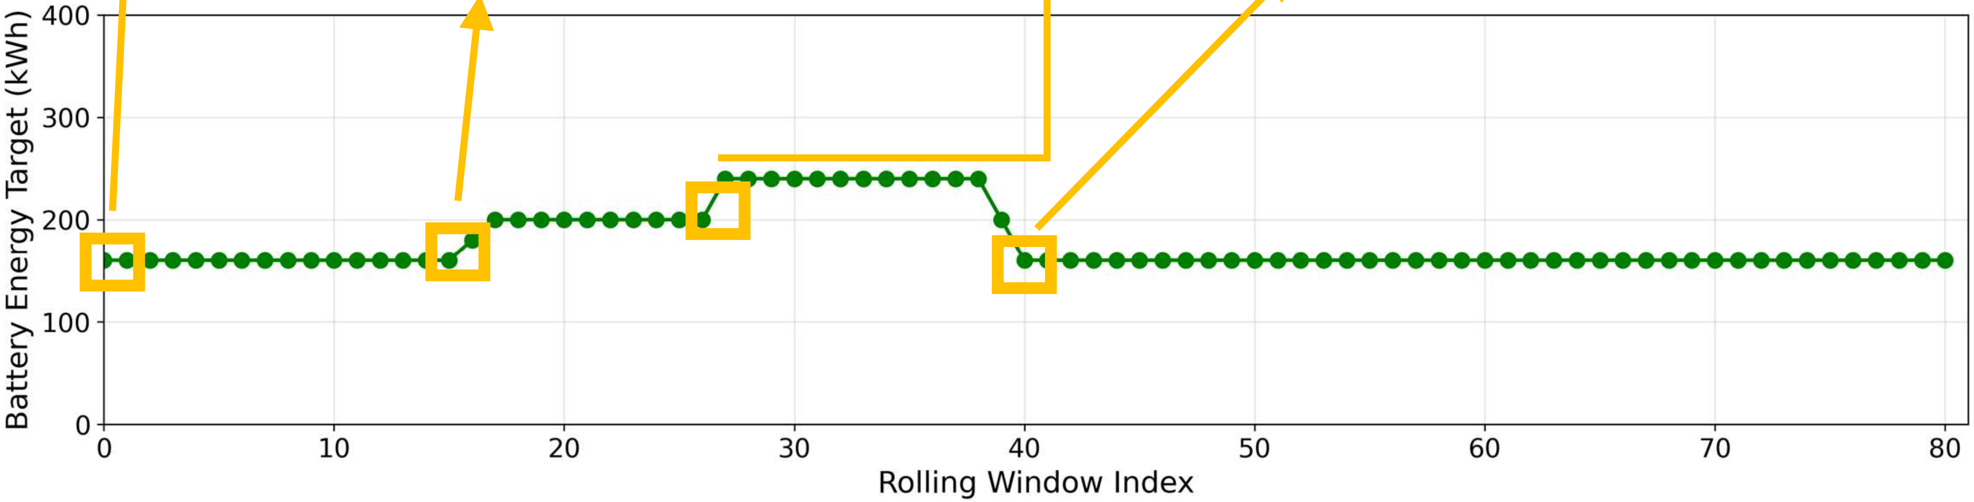

Cost Performance Summary

• **Day-Ahead Cost:** £7907
The reference day-ahead plan.

• **Test 0 (Fixed Penalty):** 7777
Baseline - Rolling optimisation with a fixed penalty yields a cost reduction compared to day-ahead.

• **Test 1 (Dynamic Penalty Adjustment):** 7777
Dynamic penalty adjustment shows similar performance to the fixed penalty approach.

• **Test 2 (SOC Constraint Adjustment):** 7718
Incorporating terminal battery SOC constraint adjustments results in the lowest cost improvement.

Concluding Remarks

- **Focus on Strategic Objectives or Constraint Management?:** Intuitively, adjusting objectives based on evolving goals will make the system operation effective. However, constraint management performs better in the shown case study.

- **AI Agents Learn and Adapt:** Leveraging techniques like supervised fine-tuning (SFT) or reinforcement learning (RL) enables AI agents to continuously improve decision-making by learning from past experience.

- **Optimisation Is Just a Tool, Not the Destination:** True energy management success lies in intelligent, adaptive **decision-making** — optimisation supports it, but does not replace it.

Thanks for your Attention!

- **My Personal Webpage:**
- https://drzekunguo.github.io/

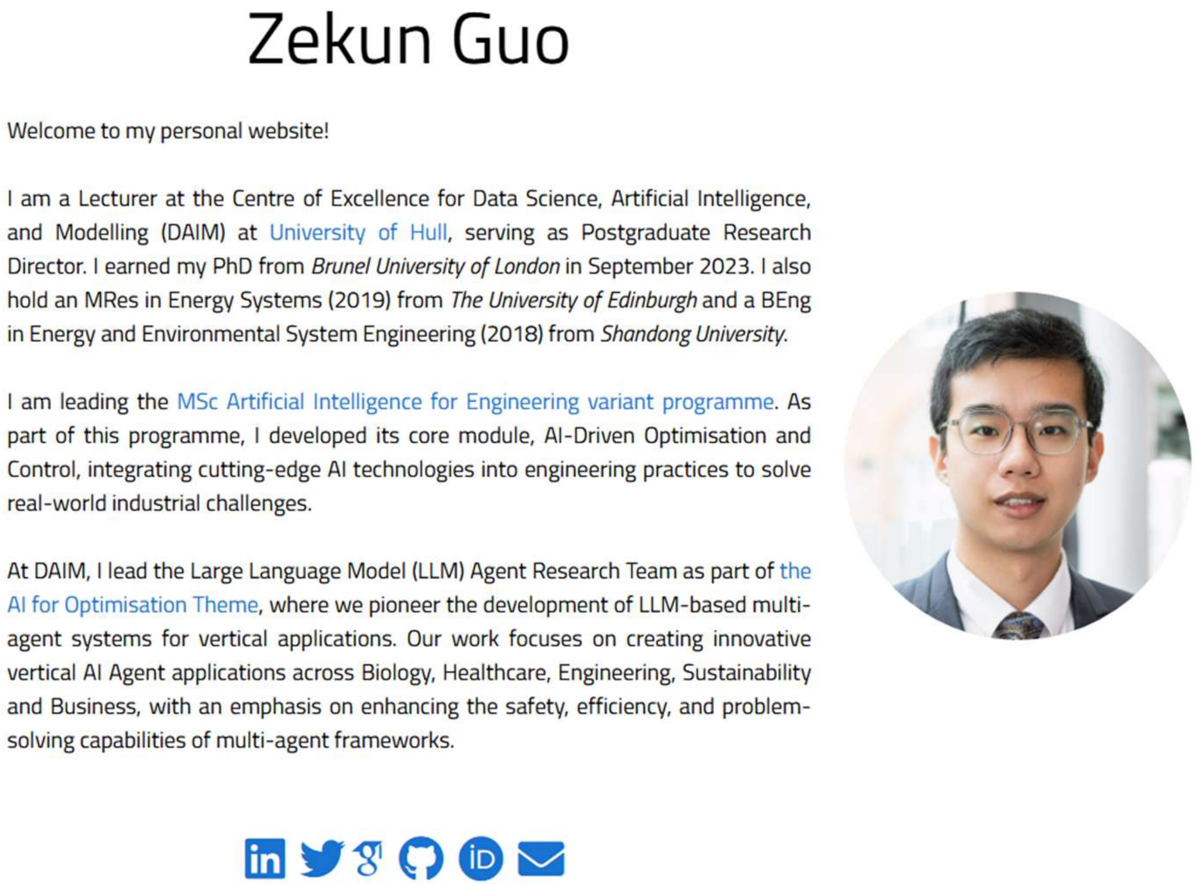

Zekun Guo

Welcome to my personal website!

I am a Lecturer at the Centre of Excellence for Data Science, Artificial Intelligence, and Modelling (DAIM) at University of Hull, serving as Postgraduate Research Director. I earned my PhD from *Brunel University of London* in September 2023. I also hold an MRes in Energy Systems (2019) from *The University of Edinburgh* and a BEng in Energy and Environmental System Engineering (2018) from *Shandong University*.

I am leading the MSc Artificial Intelligence for Engineering variant programme. As part of this programme, I developed its core module, AI-Driven Optimisation and Control, integrating cutting-edge AI technologies into engineering practices to solve real-world industrial challenges.

At DAIM, I lead the Large Language Model (LLM) Agent Research Team as part of the AI for Optimisation Theme, where we pioneer the development of LLM-based multi-agent systems for vertical applications. Our work focuses on creating innovative vertical AI Agent applications across Biology, Healthcare, Engineering, Sustainability and Business, with an emphasis on enhancing the safety, efficiency, and problem-solving capabilities of multi-agent frameworks.

References

[1] Vaswani et al (2018) **Attention is all you need**. *Proc. of NeurIPS 2018*.
[2] Yao, S et al (2023). **React: Synergizing reasoning and acting in language *models***. *ICLR 2023.*
[3] Li, G et al (2023). **CAMEL: Communicative Agents for "Mind" Exploration of Large Language Model Society**. *Proc. of NeurIPS 2023*.