

Chapter 8

Conclusion and Future Work

8.1 Conclusions

This thesis offers a brief overview of the application of ML and DL techniques in modeling and optimizing STE consumption within SBs. By utilizing various algorithms—such as LR, ARIMA, RNN, LSTM, Bi-LSTM, TCN, and GRU—the research emphasizes the creation of effective forecasting models based on multiple real-world DSs.

The goal is to establish predictive frameworks that enhance accuracy while also minimizing training time and model size, thereby making them suitable for real-time implementation in dynamic building settings. The thesis underscores the advantages of each method in tackling the intricate and non-linear characteristics of EC data, ultimately facilitating improved energy efficiency and informed decision-making in smart infrastructure.

The main findings of this research offer several important insights. Firstly, the study demonstrates the effectiveness of combining various techniques for time-series EF, highlighting how different types of input data affect prediction outcomes. It was observed that incorporating features like the current day or hour significantly

enhances the accuracy of multivariate models. Among the models assessed, TCN-Bi-LSTM showed the highest performance for univariate forecasting, while TCN-GRU proved to be the most efficient for multivariate scenarios, particularly when considering lag in hours. These results underscore the significance of thoughtful feature selection and the development of hybrid models to improve the precision of energy forecasts.

This thesis advances the state of smart-building energy analytics by proposing and rigorously evaluating a series of hybrid DL architectures. Two standouts are TCN-SA-GRU and TCN-SA-Bi-LSTM. Both begin with TCN, which excel at extracting high- and low-frequency signals from multivariate sensor streams. A self-attention layer then highlights the most relevant time steps, while the recurrent backbone—either a Gated Recurrent Unit (GRU) or a bidirectional LSTM—captures LTD that simpler models often miss. Tested on extensive real-world DSs from Berkeley and Vienna, these hybrids consistently outperform baseline methods such as ARIMA, CNN-LSTM, and TCN-RNN, delivering lower MAE and RMSE scores and higher R^2 values, particularly for 15-minute, 5-hour, and 7-hour horizons.

The research makes three key contributions. First, it shows that blending multi-scale convolutions, self-attention, and recurrent memory yields a robust framework for LTEF, even when data are noisy, irregular, or partially missing. Second, by training these networks within a federated learning setup, the study demonstrates how building owners can share model updates rather than raw data thereby maintaining privacy while benefiting from collective intelligence. Third, the integration of explainable-AI tools such as LIME uncovers which variables most notably HVAC activity drive demand, turning black-box predictions into actionable insights.

Practical implications are far-reaching. Accurate forecasts empower facility managers to schedule loads, negotiate demand-response contracts, and integrate on-site renewables more effectively. Occupants benefit through improved comfort and lower utility bills, while utilities gain a clearer view of aggregated demand, easing grid

stress. On a broader scale, reliable prediction tools support national goals of reducing greenhouse gas emissions and curbing the environmental cost of electricity generation.

The work also highlights several avenues for future research. Incorporating additional contextual factors—temperature, humidity, daylight hours, holiday calendars, or real time occupancy counts could further refine accuracy across diverse climates and building types. Lightweight, edge-based versions of the models would enable real-time control in resource-constrained IoT devices, while online learning mechanisms could adapt parameters on the fly as usage patterns evolve. Finally, expanding the DSs to include a wider array of electrical products—such as EV charging stations or advanced lighting systems will enhance model generalizability.

8.2 Future Directions

Future research should create adaptive, real-time learning systems with multi-modal data (weather, occupancy, IoT sensors, grid demand) to improve accuracy in various environments. Top priorities are: 1) Lightweight edge-compatible implementations for pre-emptive control, 2) Explainable AI methods to uncover prediction drivers, 3) Federated learning for privacy-preserving multi-building coordination, and 4) Closely integrated with building automation systems. Models must be validated across geographies and building types and include renewable energy dynamics. This double emphasis on technical sophistication (ongoing learning, edge deployment) and real-world deployment (interpretability, system integration) will drive scalable, sustainable EM solutions that reconcile efficiency, cost, and environmental objectives.