

Scaling Machine Learning for building operational phase



DAVOR STJELJA

AI Lead, Granlund Oy, Finland
PhD student, Aalto University, Finland
davor.stjelja@aalto.fi

Machine learning in building management faces challenges due to the unique characteristics of each building, hindering broad applicability and scalability. The field is evolving to address these challenges by leveraging explicit and cognitive knowledge: explicit for structured problem-solving like HVAC troubleshooting through knowledge graphs, and cognitive for understanding complex interactions through historical data analysis. Key issues include the necessity for building-specific models, ensuring model robustness, and bridging professional trust gaps. Innovative approaches like transfer learning, which involves training models on extensive data from one building and fine-tuning with data from another, and probabilistic predictions, which account for uncertainty, are showing promise. These strategies are enhancing the adaptability, scalability, and trustworthiness of machine learning solutions in building management, marking significant progress in the field.

Keywords: Scaling Machine Learning, Smart Buildings, AI for building management, knowledge graphs.

Machine Learning for effective building management

The European Union's Smart Readiness Index is at the forefront of promoting this digital shift. However, this move towards digitalization generates vast amounts of data, which can be overwhelming for manual processing.

In this scenario, cutting-edge technologies such as Artificial Intelligence (AI) and Machine Learning (ML) emerge as vital tools. They have the capability to efficiently process and interpret large data sets, playing a crucial role in the reduction of carbon emissions.

Building data is available in multiple layers of detail. At the macro level, we can track overall energy consumption. On a more detailed scale, we can examine data from specific subsystems like heating, cooling, and ventilation. Going even further, it's possible to monitor particular settings and metrics within these subsystems. Additionally, factoring in variables related to building occupants – such as their schedules, count, satisfaction, and comfort levels – alongside facility management aspects like maintenance schedules and costs, gives us a more complete picture of a building's operations. Navigating through this extensive array of data can be daunting. This is precisely where AI and ML prove to be indispensable. They enable us to make sense of and effectively utilize this vast amount of information.

Challenges of wider usage of ML

A lot of the existing research is quite localized and narrow in scope, often with researchers focusing on data from their own institutions, such as universities or campuses (Miller, 2019). In these environments, they develop and test algorithms. The major drawback of this approach is its limited external validation: these solutions aren't extensively tested across various types of buildings. This lack of diverse application testing restricts the broader implementation of machine learning in optimizing building efficiency. As a result, the potential of machine learning in this field is not fully realized, underlining the need for more generalized and widely applicable research.

The distinct characteristics of each building pose a substantial challenge to the scalability of machine learning solutions in this domain. Buildings differ significantly in aspects such as their geographic location, physical properties, technical systems, control logic, and patterns of use. For example, a building's location affects its exposure to specific climate and weather conditions. Its technical systems can vary widely in terms of age and functionality, and the control logic

implemented might be unique to that building. Moreover, even buildings that are physically identical can have vastly different usage patterns. Consequently, a machine learning model trained on the data from one building may not perform well when applied to another. This inability to generalize effectively across different buildings is a major obstacle to scaling machine learning solutions in building management.

Two knowledge types in the built environment

Addressing the challenges in scaling machine learning for buildings can be approached by harnessing the two primary types of knowledge in the built environment: explicit and cognitive, as depicted in **Figure 1**. Explicit knowledge is programmable, meaning it can be systematically codified and queried. For example, consider a scenario where a room is overly warm. You can follow a set of programmable links to pinpoint the issue: the thermometer in Room X connects to a specific ventilation system, which in turn is linked to a particular heating coil. This chain of ‘cause-and-effect’ is programmable and can be directly queried.

In contrast, cognitive knowledge emerges from observation and experience. It involves understanding complex interactions, like how room temperature varies with different heating levels, a task complicated by the multitude of influencing parameters. Machine learning excels here by learning the dynamics of a room from historical data. This process of learning and analysis helps in making accurate predictions, performing classifications, and conducting further analyses, thus enabling a deeper understanding of the more subtle and intricate aspects of building management.

Indeed, explicit knowledge in building management can be effectively captured and utilized through knowledge graphs. A knowledge graph is essentially a structured way of organizing information, enabling a computer to understand and interpret the relationships




Figure 1. Illustration of two knowledge types in the built environment: Explicit and Cognitive knowledge.

and connections between various data points or entities. In the context of building management, this becomes a powerful tool to transform intricate HVAC schematics into clear, comprehensible data structures that computers can work with.

For example, in **Figure 2** example is shown from paper by Kukkonen et al (Kukkonen et al., 2022). these graphs can illustrate the relationships between different components in a building, such as which terminals are connected to pump. This structured format of information can then be readily queried for a variety of applications. If there’s an issue with the heating system, for instance, specific prechecks can be programmed or machine learning-based queries can be utilized. The computer, using the knowledge graph much like a database, can then provide insightful answers. Knowledge graphs thus transform complex building systems into accessible and actionable data. This not only simplifies the process of troubleshooting but also enhances overall management efficiency, making the system more responsive and intelligent.

Cognitive knowledge in the context of building management is well-suited for characterization through machine learning. However, applying machine learning at scale in the built environment is fraught with challenges. Firstly, the typical approach of building unique models for each building is resource-intensive. It requires extensive manual data collection and individual model setup, leading to significant costs. Secondly, there’s the challenge of model robustness. Often, models developed for one specific context may not perform effectively when applied in a slightly different environment. This limitation can severely restrict their practical utility. Thirdly, there is a notable trust gap, especially among building professionals who may not be familiar with machine learning. The complexity and sometimes perceived opacity of these models can lead to skepticism regarding their predictions and overall reliability. Finally, the dynamic nature of buildings presents a significant hurdle. Buildings undergo continuous changes — retrofits, upgrades, or shifts in usage patterns. As a result, a model trained on historical data may quickly become outdated or inaccurate, failing to adapt to new circumstances.

Potential solutions for scaling Machine Learning

These challenges highlight the pressing need for more adaptable, robust, and transparent ML solutions in the built environment. Such solutions should not only be technically proficient but also accessible and

understandable to professionals in the field, ensuring their wider acceptance and effective integration into building management practices.

The exploration of transfer learning as a solution for scaling ML in the built environment is indeed a significant stride in this field. Transfer learning offers a practical and efficient way to apply machine learning models across different buildings, overcoming some of the key challenges associated with building-specific model development. The concept of utilizing a pre-trained model on a new building is particularly noteworthy. By training a model extensively on one building where there is a wealth of data, and then fine-tuning it with a smaller data set from a new building. This method not only saves time but also makes the process more scalable and feasible across various building types and environments.

We have worked on predicting room occupancy with transfer learning method, where I trained a model initially on a meeting room with abundant occupancy data, and then fine-tuning it with limited data from a

room in a completely different building, which showed promising results (Stjelja et al., 2022). The fact that the model performed satisfactorily despite significant differences in building size and HVAC systems is particularly promising. It underscores the potential of transfer learning to revolutionize the application of ML in the built environment, making it more scalable, efficient, and adaptable to varying building dynamics. This approach could indeed be a game-changer in enhancing the application of ML for building management and energy efficiency.

Another promising avenue for scaling ML in building applications is the use of probabilistic predictions, a method we explore in depth in my research (Stjelja et al., n.d.). This approach enhances model robustness and bolsters user confidence in AI systems. Adopting probabilistic predictions moves beyond the limitations of single-point forecasting by embracing a methodology that inherently accounts for uncertainty. The concept of estimating an entire distribution, rather than a single outcome, brings a critical dimension of realism and practicality to ML models. By incorporating uncertainty

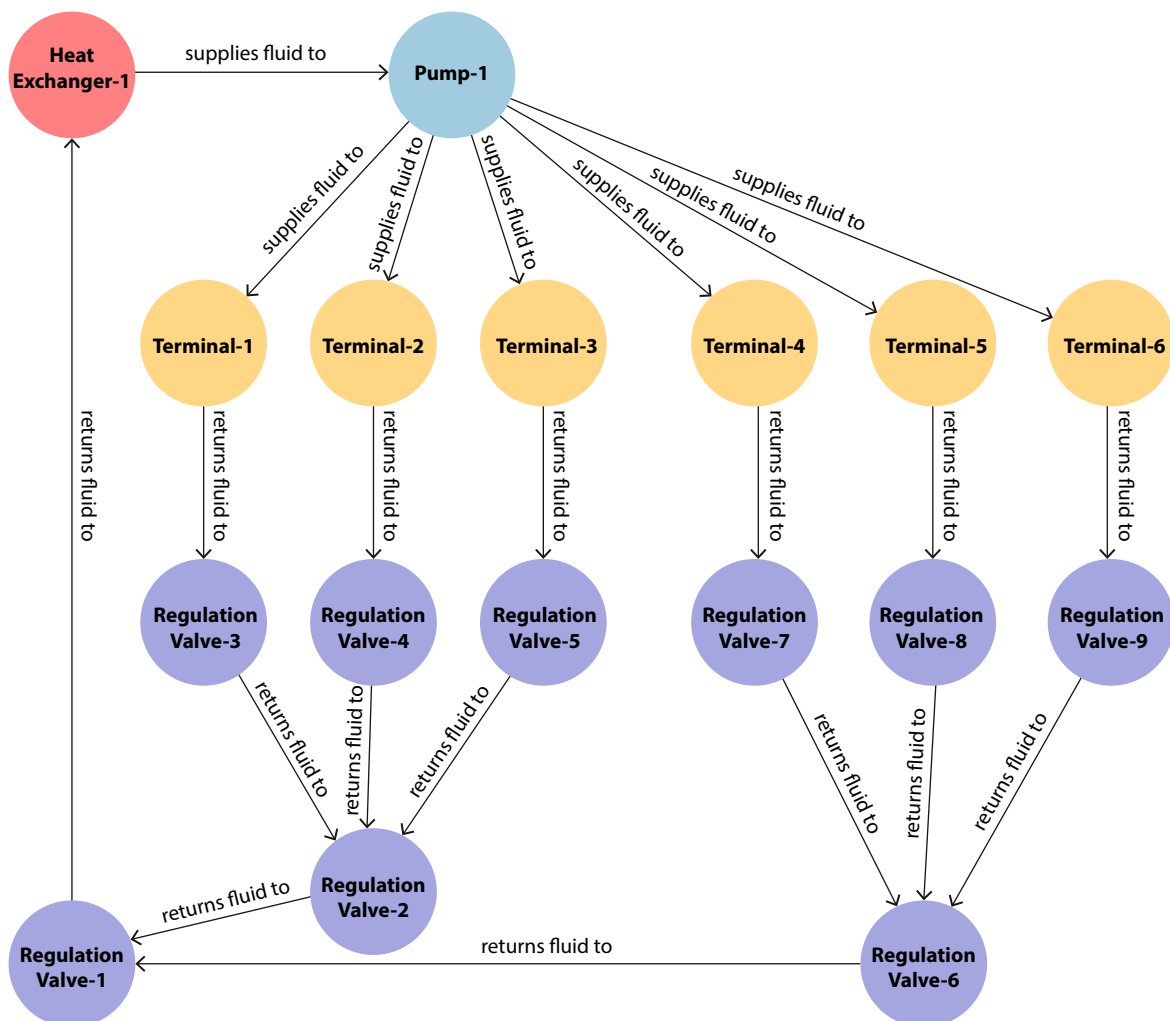


Figure 2. Example of knowledge graph showing fluid flow from heat exchanger through a pump to radiators. Kukkonen et al.

quantification, these models not only provide predictions but also communicate the confidence level or potential variability in these predictions, potentially bridging the gap in trust and confidence among professionals skeptical of ML's black-box nature. This aspect is particularly valuable in complex systems like buildings, where numerous variables and unpredictable factors come into play. In this paper we compare two approaches of predicting building energy consumption between two probabilistic algorithms. Furthermore, these predictions are then used for the detection of drift anomalies. The proposed method doesn't alert to immediate issues but alerts of emerging trends or irregularities that could become problematic if unaddressed.

Promising directions

The identified areas of research present promising directions for future investigations aimed at enhancing the scalability of machine learning for building operations. Each area offers distinct solutions and perspectives that could considerably propel the field forward.

- **Continual Learning:** This area focuses on the dynamic aspect of machine learning models, particularly their ability to adapt continually as buildings evolve. This is crucial in the context of buildings, which undergo regular changes in structure, usage, or systems. Continual learning ensures that machine learning models remain relevant and accurate over time, adjusting to new data and conditions. Example is this large-scale comparison and demonstration of continual learning in building operation (Li et al., 2023).
- **Explainable AI (XAI):** The field of Explainable AI is gaining traction, particularly for its potential to demystify the decision-making processes of AI models. By making AI predictions more transparent and understandable, XAI could greatly enhance trust and confidence among building management professionals. This is particularly important for those who might be hesitant to rely on AI due to its perceived opacity. This review paper shows state of the art research on XAI topic in this field (Chen et al., 2023).
- **Few-Shot Learning:** This approach is particularly relevant for situations where data is scarce, such as in new or recently retrofitted buildings. Few-shot learning allows for the training of effective machine learning models using minimal data points, which is a significant advantage in scenarios where extensive historical data is not available. Interesting paper using few-shot building energy prediction is (Tang et al., 2023).
- **Transfer Learning:** For a comprehensive understanding of transfer learning in building management, review paper by Pinto et al. could be a valuable resource (Pinto et al., 2022). Transfer learning, as highlighted, can address the issues of scalability and data requirement in deploying machine learning models across different buildings.

These research areas collectively represent the forefront of AI and ML in building management. They offer promising solutions to overcome current challenges and pave the way for more efficient, adaptable, and trustworthy AI applications in this domain. For professionals and researchers interested in this field, delving into these topics would provide a deeper understanding of the potential and direction of AI in the built environment. ■

References

- Chen, Z., Xiao, F., Guo, F., & Yan, J. (2023). Interpretable machine learning for building energy management: A state-of-the-art review. *Advances in Applied Energy*, 9, 100123. <https://doi.org/10.1016/j.adapen.2023.100123>.
- Kukkonen, V., Kückükavci, A., Seidenschur, M., Rasmussen, M. H., Smith, K. M., & Hviid, C. A. (2022). An ontology to support flow system descriptions from design to operation of buildings. *Automation in Construction*, 134, 104067. <https://doi.org/10.1016/j.autcon.2021.104067>.
- Li, A., Zhang, C., Xiao, F., Fan, C., Deng, Y., & Wang, D. (2023). Large-scale comparison and demonstration of continual learning for adaptive data-driven building energy prediction. *Applied Energy*, 347, 121481. <https://doi.org/10.1016/j.apenergy.2023.121481>.
- Miller, C. (2019). More Buildings Make More Generalizable Models—Benchmarking Prediction Methods on Open Electrical Meter Data. *Machine Learning and Knowledge Extraction*, 1(3), Article 3. <https://doi.org/10.3390/make1030056>.
- Pinto, G., Wang, Z., Roy, A., Hong, T., & Capozzoli, A. (2022). Transfer learning for smart buildings: A critical review of algorithms, applications, and future perspectives. *Advances in Applied Energy*, 5, 100084. <https://doi.org/10.1016/j.adapen.2022.100084>.
- Stjelja, D., Jokisalo, J., & Kosonen, R. (2022). Scalable Room Occupancy Prediction with Deep Transfer Learning Using Indoor Climate Sensor. *Energies*, 15(6), 2078. <https://doi.org/10.3390/en15062078>.
- Stjelja, D., Kuzmanovski, V., Kosonen, R., & Jokisalo, J. (n.d.). Building consumption anomaly detection: A comparative study of two probabilistic approaches. *Submitted to Energy & Buildings*.
- Tang, L., Xie, H., Wang, X., & Bie, Z. (2023). Privacy-preserving knowledge sharing for few-shot building energy prediction: A federated learning approach. *Applied Energy*, 337, 120860. <https://doi.org/10.1016/j.apenergy.2023.120860>.