

Towards data-driven simulation models for building energy management

Juan Gómez-Romero¹[0000-0003-0439-3692] and
Miguel Molina-Solana¹[0000-0001-5688-2039]

Dept. Computer Science and AI, Universidad de Granada, Spain
{jgomez,miguelmolina}@decsai.ugr.es

Abstract. The computational simulation of physical phenomena is a highly complex and expensive process. Traditional simulation models, based on equations describing the behavior of the system, do not allow generating data in sufficient quantity and speed to predict its evolution and make decisions accordingly automatically. These features are particularly relevant in building energy simulations. In this work, we introduce the idea of deep data-driven simulation models (D3S), a novel approach in terms of the combination of models. A D3S is capable of emulating the behavior of a system in a similar way to simulators based on physical principles but requiring less effort in its construction—it is learned automatically from historical data—and less time to run—no need to solve complex equations.

Keywords: Data-driven simulation model · Deep Learning · Building Energy Management.

1 Introduction

According to a 2019 report by the consulting firm ABI Research [14], in the next five years, it is expected that more than 100,000 companies around the world will use simulation software, implying a business volume of over 2,500 million dollars annually in 2025.

However, computer simulation of physical phenomena, such as meteorology, energy transfer, or nuclear reactions, is costly. On the one hand, to create a simulation model of a system, it is necessary that the relationship between inputs and outputs is known and can be expressed in a calculable way. Thus, these models are generally created manually by coding the equations that describe the physical behavior of the system. On the other hand, running a complex simulation model can take several hours (or even days) and require large amounts of computational resources. Consequently, in most common problems, it is impossible to use these models to predict the evolution of the system in real-time and automatically make decisions from these simulations.

Taking the energy behavior of buildings (residential and non-residential) as an example, a physical simulation model characterizes the response of its components to the action of the equipment and the environmental conditions, employing differential equations that reproduce the energy transfer laws. These

models are built with specialized applications (e.g., Modelica or EnergyPlus) by assembling predefined modules that imitate the thermal response of different structures, materials, and equipment. Leveraging a simulation model, one can build an automatic control software that estimates the behavior of the building under different operating sequences (air conditioning, lights, etc.) and selects the one that involves the lowest energy consumption while maintaining comfort conditions (indoor temperature, humidity, CO₂ concentrations, etc.).

This approach is known as Model Predictive Control (MPC) [11] and can be applied in many areas beyond energy. Solutions based on MPC offer numerous advantages over traditional reactive controllers, especially in terms of optimizing the process in the medium / long term by taking into account the inertia of the systems. However, its implementation is limited for the problems mentioned above: creating the models requires much human effort, and their execution takes too long to generate operation plans within a reasonable time dynamically [9].

Although the literature has been demonstrated the possibility of significantly improving the control of a building and reducing energy consumption [7], numerous difficulties were also encountered in extending the approach to other contexts. The main bottleneck is to develop the physical simulation models that the control algorithm uses since these are created from scratch and cannot be reused from one building to another. This problem is even more acute when trying to apply Deep Reinforcement Learning techniques in the field of energy control [6]. In these cases, traditional simulation models cannot generate data in sufficient quantity and speed to train and validate the proposed algorithms.

As an alternative to physical simulation models, some approaches have been proposed in the literature to create prediction models of the behavior of systems using Machine Learning. However, these prediction models alone are not capable of addressing various needs that a simulation model must satisfy, such as the stability of the model against minor variations in inputs, the influence of the environment on the behavior of the system, the possibility of modifying its behavior through control instructions, or the use of sensor data affected by imprecision and uncertainty.

For these reasons, we aim to develop new algorithms, based on Deep Learning, to automatically learn fast, accurate, and realistic simulation models of a physical system from data. These models could be used in numerous applications and particularly in the MPC processes for energy optimization mentioned above.

The rest of this paper summarizes the background and state of the art (Section 2), and describes the key concepts and approach of our proposal (Section 3).

2 Background

Various proposals in the literature aim to perform computational simulation of physical systems using numerical models learned from available data. This approach is a generalization of the concept of system identification, a discipline close to classical control that studies the calculation of the parameters of a predefined model to adjust its outputs to those of the real system [13]. Traditionally,

works on system identification have used algebraic and statistical methods [5], including time series analysis by autoregression with ARIMA or ARIMAX [19]. This is because the knowledge and availability of these tools have been traditionally broader, although in most cases, they are not the most effective, as we recently concluded in a review of works in the field of energy efficiency [16].

It was not until recently that machine learning techniques began to be applied in system identification [3]. For example, in [1] techniques based on Gaussian kernels are used to emulate the behavior of molecules at the electronic level without the need to solve differential equations. In contrast, in [10] similar techniques are applied to recognize galaxies from spectral data analysis, a process that typically requires running multiple simulations. In both cases, the proposed solutions manage to approximate the systems accurately because 1) there is a reduced number of output variables, and 2) aspects of the problem are encoded in the own model (e.g., which variables are relevant, what structure have the kernels that define the process). On the other hand, it is not easy to extend these models to other settings, even if they are only slightly different.

Deep Learning techniques have been investigated to learn [17] and calibrate [22] data series prediction models to address these limitations in recent years. They allow the automatic extraction of system characteristics and achieve more precision in the data series results than classical techniques. Among the many possible architectures, recurrent neural networks (RNNs), which allow cycles in the calculation graph, are the most effective for modeling the temporal behavior of a dynamic system [2], improving the results of the classic autoregressive techniques [18]. On the other hand, convolutional neural networks (CNNs), which are used mainly for image processing, allow ordered data sequences to be processed [8], although their adjustment is complex even for simple problems.

Simulation using Deep Learning techniques is, therefore, a new area of research in which there are hardly any works that exploit the capabilities of deep neural networks for processing multivariate data series from sensors. Furthermore, due to their own characteristics, neural networks present various additional problems regarding the connection of the prediction models with the actual phenomenon: robustness in the face of variations in the inputs, detection of errors in the data, characterization of cause-effect relationships, etc. Thus, for example, in the field of energy efficiency of buildings, there are some results about the modeling of thermal behavior [4,12], but on a very small scale, with a very short time horizon, and with little explanatory capacity.

In contrast, RNNs have been very successful in the field of natural language processing, as they are capable of learning predictive models of language. These networks are widely used in machine translation so that the network obtains the phrase in the target language that most closely matches the phrase in the source language. This type of architecture, called sequence-to-sequence [15], implements a procedure that first encodes the input phrase as a sequence of numbers (embeddings) and then decodes them to form the output phrase. Various improvements to this architecture, such as transformers [20], are close to the precision of a human translator in non-specialized domains.

3 Proposal

The fundamental concept we based our proposal on is that of a data-driven simulation model (DDS). A DDS model is capable of emulating the behavior of a system in a similar way to that of traditional simulators based on physical principles, but requiring less effort in its construction—it is automatically learned from historical data—and less time for its execution—no need to perform complex calculations. Recent advances in the area of Deep Learning suggest that it is possible to create DDS models based on deep neural networks¹, which we will call D3S (deep data-driven simulation models), improving the prediction capabilities of current time series algorithms. The D3S concept is very close to that of the digital twin [21], highlighting the essential properties of creation from data and the use of neural networks. Although neural networks are usually highly costly in terms of training, inference (simulation, in our use case) can be performed efficiently and quickly.

Formulation: A general simulation model can be seen as a computational process that transforms several inputs—corresponding to the previous states of the system, the applied control signals, and the expected conditions of the environment from the current instant—into an output that represents the sequence of n states through which the system passes from the current state to a final state (Figure 1). Usually, the environmental conditions are not known a priori and can be estimated or even unknown. As for the control signals, they can be the result of an automatic optimization process. In the simplest case, the simulator would only consider a previous state and a following state. Likewise, it may happen that the system is not controllable or that the environmental conditions are not relevant. In more complex cases, in addition to the previous states, it would be necessary to incorporate the environmental conditions and the previous control actions. For simplicity, and without loss of generality, we will keep the general formulation of Figure 1.

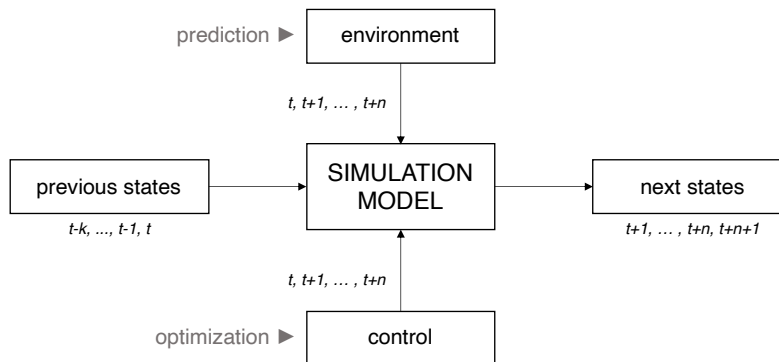


Fig. 1. Schematic view of a general simulation model.

¹ see, for instance, ICLR 2021’s workshop “Deep Learning for Simulation (SIMDL)”

Approximation: We formalize the data-driven simulation problem as a multivariate time series prediction problem. Thus, the simulation model learns to predict a series of data representing the following states from the series of data representing the previous states, the environment, and the control, each of them including observations of several variables. As explained, classical data series analysis techniques are insufficient to solve this problem since they present difficulties in predicting more than one variable at the output. They also have limitations when it comes to capturing non-linear relationships between inputs and outputs. For these reasons, there is a need for new techniques in Deep Learning for handling data series (Figure 2). A source of inspiration is automatic translation, which obtains the sequence of words in the target language that best represents the sequence of words in the original language, taking into account the context and the lexical-semantic relationships between them. Similarly, we propose the creation of algorithms capable of transforming the sequences of input values (environment, previous states, and control actions) into sequences of values (following states) with a relationship that is not necessarily linear considering the interrelation between states, control, and environment.




Fig. 2. Schematic view of our proposal for a Data-driven deep simulation model (D3S).

Challenges: This approach to data-driven simulation faces several challenges. In the following, we discuss several aspects that must be taken into account in building energy management:

1. Data availability: Data-driven simulations require a considerable amount and diversity of data to be performed. While data is generally available in modern buildings from SCADA systems, it is less commonly representative of exceptional or anomalous situations. Therefore, the scale of data needed by modern Deep Learning techniques, particularly transformer architectures, may exceed what can be obtained from a regular building.
2. Physics constraints: Physical systems are subject to bounding conditions, which should be incorporated into a D3S model. Neural networks can incorporate such conditions as regularizations, i.e., penalties in loss functions.

However, it is not trivial to translate from one language to another, i.e., PDEs to regularization terms. Hence, there is a growing interest in physics-informed machine learning, which investigates new neural architectures that integrate PDEs as additional optimization targets.

3. Dynamic behavior: Simulation models must evolve as their underlying process. This fact is commonplace in buildings, subject to renovations, aging, and changes in uses. Consequently, detecting model shift and activating recalibration must be automatically performed. Studies on continuous learning and data assimilation with Deep Learning suggest that this goal is achievable, but it may also impact the models' stability.
4. Explainability: Building management systems are cyber-physical systems that involve human operators' participation and sometimes the realization of critical tasks (e.g., CO_2 control). Therefore, the use of D3S for automated (or semi-automatic) decision-making requires at least some explanation of the outputs. Besides, experts' experience may be beneficial to bootstrap the training of D3S (e.g., by pre-identified relevant variables, periodicity, or spurious correlations). The growing body of knowledge on explainable AI should play a role here to make D3S more interpretable.
5. Computational cost: Learning a D3S model remains an expensive process that may require substantial computational resources, time, and energy. Accordingly, measuring the environmental impact of these models is essential to evaluate energy savings precisely. Model reuse from one building to another by transfer learning followed by fine-tuning could significantly reduce the data and energy needed to train the models.

References

1. Brockherde, F., Vogt, L., Li, L., Tuckerman, M.E., Burke, K., Müller, K.R.: Bypassing the kohn-sham equations with machine learning. *Nature Communications* **8**, 1–10 (2017). <https://doi.org/10.1038/s41467-017-00839-3>
2. Che, Z., Purushotham, S., Cho, K., Sontag, D., Liu, Y.: Recurrent neural networks for multivariate time series with missing values. *Sci Rep* **8**, 1–12 (2018). <https://doi.org/10.1038/s41598-018-24271-9>
3. Chiuso, A., Pillonetto, G.: System identification: A machine learning perspective. *Annual Review of Control, Robotics, and Autonomous Systems* **2**, 281–304 (2019). <https://doi.org/10.1146/annurev-control-053018-023744>
4. Ferracuti, F., Fonti, A., Ciabattoni, L., Pizzuti, S., Arteconi, A., Helsen, L., Comodi, G.: Data-driven models for short-term thermal behaviour prediction in real buildings. *Applied Energy* **204**, 1375–1387 (2017). <https://doi.org/10.1016/j.apenergy.2017.05.015>
5. Gevers, M.: A personal view of the development of system identification: A 30-year journey through an exciting field. *IEEE Control Systems* **26**, 93–105 (2006). <https://doi.org/10.1109/MCS.2006.252834>
6. Gómez, J., Molina-Solana, M.: Towards self-adaptive building energy control in smart grids. In: *NeurIPS 2019 Workshop Tackling Climate Change with Machine Learning*. Vancouver, Canada (Dec 2019), <https://www.climatechange.ai/papers/neurips2019/49>

7. Gómez-Romero, J., Fernández-Basso, C.J., Cambronero, M.V., Molina-Solana, M., Campaña, J., Ruiz, M.D., Martin-Bautista, M.J.: A probabilistic algorithm for predictive control with full-complexity models in non-residential buildings. *IEEE Access* **7**, 38748–38765 (2019). <https://doi.org/10.1109/ACCESS.2019.2906311>
8. Kasim, M.F., Watson-Parris, D., Deaconu, L., Oliver, S., Hatfield, P., Froula, D.H., Gregori, G., Jarvis, M., Khatiwala, S., Korenaga, J., Topp-Mugglestone, J., Viezzer, E., Vinko, S.M.: Building high accuracy emulators for scientific simulations with deep neural architecture search (2020)
9. Killian, M., Kozek, M.: Ten questions concerning model predictive control for energy efficient buildings. *Building and Environment* **105**, 403–412 (2016). <https://doi.org/10.1016/j.buildenv.2016.05.034>
10. Kwan, J., Heitmann, K., Habib, S., Padmanabhan, N., Lawrence, E., Finkel, H., Frontiere, N., Pope, A.: Cosmic emulation: Fast predictions for the galaxy power spectrum. *The Astrophysical Journal* **810** (2015). <https://doi.org/10.1088/0004-637X/810/1/35>
11. Lee, J.H.: Model predictive control: Review of the three decades of development. *International Journal of Control, Automation and Systems* **9** (2011). <https://doi.org/10.1007/s12555-011-0300-6>
12. Liu, Y., Dinh, N., Sato, Y., Niceno, B.: Data-driven modeling for boiling heat transfer: Using deep neural networks and high-fidelity simulation results. *Applied Thermal Engineering* **144**, 305–320 (2018). <https://doi.org/10.1016/j.applthermaleng.2018.08.041>
13. Ljung, L.: *System identification: theory for the user*. Prentice Hall (1999)
14. Loten, A.: More manufacturers bet on simulation software. *The Wall Street Journal* (Feb 2020), <https://www.wsj.com/articles/more-manufacturers-bet-on-simulation-software-11582240105>
15. Luong, M.T., Le, Q.V., Sutskever, I., Vinyals, O., Kaiser, L.: Multi-task sequence to sequence learning. In: *International Conference on Learning Representations (ICLR 2016)* (2016), <http://arxiv.org/abs/1511.06114>
16. Molina-Solana, M., Ros, M., Ruiz, M.D., Gómez-Romero, J., Martin-Bautista, M.J.: Data science for building energy management: a review. *Renewable and Sustainable Energy Reviews* **70**, 598–609 (2017). <https://doi.org/10.1016/j.rser.2016.11.132>
17. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* **378**, 686–707 (2019). <https://doi.org/10.1016/j.jcp.2018.10.045>
18. Shi, H., Xu, M., Li, R.: Deep Learning for Household Load Forecasting—A Novel Pooling Deep RNN. *IEEE Transactions on Smart Grid* **9**, 5271–5280 (2018). <https://doi.org/10.1109/TSG.2017.2686012>
19. Shumway, R.H.: *Time series analysis and its applications*. Springer (2017)
20. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Proceedings of the 31st Conference on Neural Information Processing Systems*. Long Beach, CA, USA (Dec 2017)
21. Wagg, D.J., Worden, K., Barthorpe, R.J., Gardner, P.: Digital Twins: State-of-the-Art and Future Directions for Modeling and Simulation in Engineering Dynamics Applications. *ASME J. Risk Uncertainty Part B* **6**(3) (2020). <https://doi.org/10.1115/1.4046739>
22. Zhu, J., Hu, S., Arcucci, R., Xu, C., Zhu, J., ke Guo, Y.: Model error correction in data assimilation by integrating neural networks. *Big Data Mining and Analytics* **2**, 83–91 (2019). <https://doi.org/10.26599/BDMA.2018.9020033>