Physics-Informed Deep Operator Networks for Real-Time Spatiotemporal Monitoring of Indoor Air Quality

Minh T.N Nguyen [1 2] **Hieu V. Vu** [2 3] **Nguyen Dinh Phong** [1] **Nhan Dinh Ngo** [2] **Hai Hung Pham** [1 4] **Dung D. Le** [1 2 4 5]

[1]College of Engineering & Computer Science, VinUniversity, Vietnam
[2]VinUni-Illinois Smart Health Center, VinUniversity, Vietnam
[3]CFDWAYS Engineering Limited Liability Company, Hanoi, Vietnam
[4]Center of Environmental Intelligence, VinUniversity, Vietnam
[5]Center for AI Research, VinUniversity, Vietnam

Corresponding author email(s): hung.ph@vinuni.edu.vn, dung.ld@vinuni.edu.vn

1. Introduction

Indoor Air Quality (IAQ) is a primary determinant of occupant health, cognitive performance, and the transmission of airborne pathogens [1, 2, 3]. Klepeis et al. [4] showed that individuals spend approximately 80% to 90% of their time within built environments. Gao et al. [5] also noted that surveys indicate people spend at least 80% of their time indoors, citing the foundational work of Klepeis et al. (2001). Given this significant exposure, maintaining optimal IAQ is a global health imperative, particularly as air pollution is linked to roughly seven million deaths annually [6]. Furthermore, built environments are significant drivers of global resource use, with HVAC systems accounting for nearly 40% of global energy consumption [1, 7].

Current IAQ monitoring typically relies on sparse, wall-mounted sensors or urban air quality stations. However, these physical networks face significant economic and technical hurdles; urban monitoring stations can cost 200,000 USD to construct and 30,000 USD annually to maintain [6]. In practical applications, the limited number of sensors often results in “blind spots,” failing to detect localized pollutant hotspots or stagnant “dead zones” [8, 9]. Moreover, most existing control systems operate under a “well-mixed” assumption, treating a room as a single uniform node [5, 10].

In reality, indoor airflow exhibits significant spatiotemporal heterogeneity, meaning a single sensor reading rarely represents the pollutant concentration at an occupant’s breathing level [1, 10]. Effective IAQ management requires a transition to frameworks capable of reconstructing the full 3-dimensional (3D) pollutant field in real-time. High-fidelity modeling is essential to prevent “over-ventilation” (which wastes energy) and “under-ventilation” (which compromises safety), allowing HVAC systems to adapt to dynamic occupancy and shifting pollutant distributions [1, 8].

2. Proposed Framework

Computational Fluid Dynamics (CFD) is the standard tool for capturing detailed 3D airflow, but its high computational cost—requiring hours to simulate mere minutes of airflow—limits its use in real-time control [1, 5]. To bridge this gap, research has shifted toward data-driven surrogate models, including Artificial Neural Networks (ANN), Convolutional Neural Networks (CNN), and Graph Neural Networks (GNN) [2, 10]. Recent advancements in operator learning, such as Deep Operator Networks (DeepONet) and Fourier Neural Operators (FNO), have introduced the ability to learn mappings between infinite-dimensional function spaces, enabling resolution-invariant predictions at arbitrary spatial coordinates [11, 5, 1].

Despite the speed of data-driven models, most existing surrogates operate as “black boxes” that lack physical consistency, often failing to obey fundamental conservation laws such as mass and momentum [12, 3]. In addition, purely virtual reconstructions can deviate from empirical reality due to unresolved sub-grid dynamics. There is a need for a hybrid layer using physical sensors to provide real-time corrective values. Furthermore, the calibration of virtual sensing must require minimum effort across varying contexts, such as rooms with or without complex interiors and obstructions. This will benefit the calibration of virtual sensors as changes in interior configurations or occupants without the need for thousands of new simulations. Thus, there remains a critical need for a framework that combines CFD-level physical accuracy with the rapid inference speeds required for real-time building automation.

This study proposes a hybrid monitoring framework based on a Physics-Informed Deep Operator Network (PI-DeepONet) for the real-time spatiotemporal reconstruction of 3D IAQ [1, 5]. By embedding governing physical laws (Navier-Stokes and advection-diffusion equations) directly into the model’s loss function, the framework ensures physical consistency even with sparse data [12, 1]. The goal is to implement a hybrid concept where real-time physical sensor data is used to dynamically correct the outputs of the PI-DeepONet. The “neighbor aggregation” mechanism allows the model to treat the physics-informed prediction as a global baseline while using local sensor readings to refine high-frequency variations and resolve stochastic disturbances [6, 13].

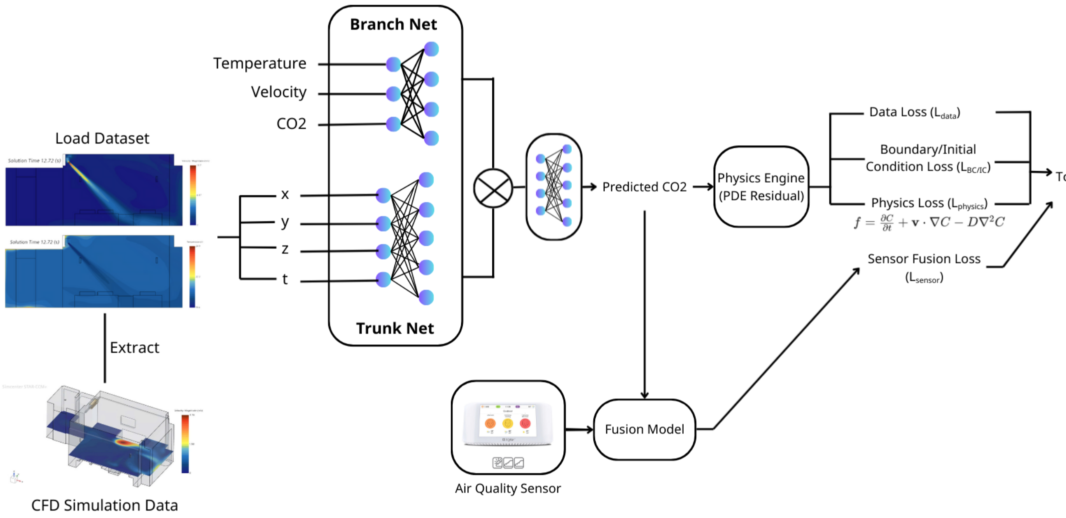

Fig. 1: Schematic of the proposed PI-DeepONet framework. The workflow proceeds from left to right: (1) **Data Generation** (left) extracts high-fidelity fields from CFD simulations. (2) The **DeepONet Architecture** (center) processes inputs via two parallel networks: a *Branch Net* encoding flow variables (Temperature, Velocity, CO_2) and a *Trunk Net* encoding spatiotemporal coordinates (x, y, z, t). These are combined to generate the **Predicted CO_2**. (3) The **Physics Engine** (right) enforces physical laws by computing the advection-diffusion PDE residual. The model is trained by minimizing a composite **Total Loss** ($\mathcal{L}_{total}$) comprising data, boundary, sensor fusion, and physics losses, ensuring the output respects both empirical data and mass conservation laws.

3. Methodology

The proposed framework integrates high-fidelity CFD data, operator learning, and sparse sensor observations into a unified real-time monitoring system. The workflow consists of three core components: (1) Data generation via CFD, (2) Offline training of the PI-DeepONet, and (3) Real-time correction using a sensor fusion module.

3.1 Case Description and CFD Dataset Settings

A realistic three-dimensional indoor domain representing a hotel guest room was modeled, explicitly resolving HVAC supply and return vents and interior furniture as flow obstructions. Transient CFD simulations were conducted to capture the unsteady evolution of airflow and thermal fields and to generate a high-fidelity reference dataset for PI-DeepONet model development. The simulations employed a URANS framework to balance physical fidelity and computational tractability for parametric studies. Numerical verification was ensured through mesh-independence assessment, while validation was conducted by comparing predicted indoor variables against measurements obtained from an iAQ Sensor Pro (AirVisual). A structured Design of Experiments was defined by systematically varying supply air temperature setpoints, inlet airflow rates, occupants and interior configurations with and without furniture to ensure coverage of representative indoor operating conditions.

3.2 Sensor Fusion Mechanism

While PI-DeepONet provides a physically robust global baseline, it may deviate due to unmodeled stochastic disturbances. To bridge the gap between simulation and reality, we employ a “neighbor aggregation” fusion mechanism. Real-time data from sparse physical sensors are compared with the model’s predictions at corresponding locations. The calculated residuals are then interpolated locally to correct the predicted field in the vicinity of the sensors, allowing the system to capture dynamic hotspots that the offline model might miss while maintaining global physical validity.

4. Conclusion

This study introduces a novel PI-DeepONet framework for real-time, 3D IAQ monitoring. By hybridizing operator learning with physical constraints and sensor fusion, the system overcomes the computational bottleneck of CFD and the accuracy limitations of pure data-driven models. Future work will focus on deploying the framework in a live building management system (BMS) to demonstrate closed-loop energy optimization.

Acknowledgments

This work is financially supported by VinUniversity under Grant No VUNI.GT.NTKH.03.

References

[1] Y. Bian, O. Schmidt, and Y. Shi. Operator learning for energy-efficient building ventilation control with computational fluid dynamics simulation of a real-world classroom. *arXiv preprint arXiv:2504.21243v2*, 2025.

[2] H. Y. Kek, A. Bazgir, H. Tan, C. T. Lee, T. Hong, M. H. D. Othman, Y. V. Fan, M. N. H. Mat, Y. Zhang, and K. Y. Wong. Particle dispersion for indoor air quality control considering air change approach: A novel accelerated cfd-dnn prediction. *Energy & Buildings*, 306:113938, 2024.

[3] I. Reda, M. B. Soultanzadeh, A. Taha, K. Zheng, D. Qi, M. M. Ouf, and L. L. Wang. Rapid indoor airflow prediction using a hybrid residual learning regression model. *Energy & Buildings*, 352:116827, 2026.

[4] Neil E Klepeis, William C Nelson, Wayne R Ott, John P Robinson, Andy M Tsang, Paul Switzer, Joseph V Behar, Stephen C Hern, and William H Engelmann. The national human activity pattern survey (nhaps): a resource for assessing exposure to environmental pollutants. *Journal of Exposure Science & Environmental Epidemiology*, 11(3):231–252, 2001.

[5] H. Gao, W. Qian, J. Dong, and J. Liu. Rapid prediction of indoor airflow field using operator neural network with small dataset. *Building and Environment*, 251:111175, 2024.

[6] Y. Feng, Q. Wang, Y. Xia, J. Huang, S. Zhong, and Y. Liang. Spatio-temporal field neural networks for air quality inference. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, 2024.

[7] Q. Zhou and R. Ooka. Implementation of a coupled simulation framework with neural network and modelica for fast building energy simulation considering non-uniform indoor environment. *Building and Environment*, 211:108740, 2022.

[8] Y. Li, Z. Tong, D. Westerdahl, and S. Tong. Predictive monitoring of built thermal environment using limited sensor data: A deep learning-based spatiotemporal method. *Sustainable Energy Technologies and Assessments*, 66:103823, 2024.

[9] M. Moradi et al. Placeholder title for moradi et al 2026 based on context. *Journal Placeholder*, 2026. Citation provided in text without full reference details in list.

[10] W. Zhang, Z. Wang, Y. Liu, C. Zhou, C. Zhang, Y. Zhao, and Y. Hu. Convolutional neural networks-based surrogate model for fast computational fluid dynamics simulations of indoor airflow distribution. *Energy & Buildings*, 326:115020, 2025.

[11] S. Yang, Y. Lee, and N. Kang. Data-efficient deep operator network for unsteady flow: A multi-fidelity approach with physics-guided subsampling. *arXiv preprint arXiv:2503.17941v2*, 2025.

[12] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 2021.

[13] Z. Wang and R. Han. Deep learning for 3d reconstruction and trajectory prediction of dust and polluted aerosols in educational environments. *Frontiers in Environmental Science*, 13:1582806, 2025.