Operator Learning for Indoor Air Dynamics: Modeling and Applications

Yuexin Bian [1] **Yuanyuan Shi** [1]

Abstract

Indoor air quality plays a critical role in ensuring occupant health, comfort, and energy-efficient building operation. Accurate prediction of indoor airflow and precise boundary-level HVAC control can significantly enhance building performance. While Computational Fluid Dynamics (CFD) offers high-fidelity modeling, its computational cost makes it impractical for real-time applications. To address this, we propose an ensemble neural operator transformer (ENOT) that predicts the spatiotemporal evolution of indoor CO_2 levels, achieving a 250,000× speed-up over traditional CFD simulations. Our contributions include a high-fidelity CFD-based dataset, a simulation pipeline for realistic indoor air modeling, and an ensemble neural operator learning framework for accurate, real-time inference. We further outline future directions in data-driven model-based HVAC control, bridging the gap between high-fidelity simulation and intelligent building management. Our code and data are publicly available at https://huggingface.co/datasets/alwaysbyx/Bear-CFD-dataset.

1. Introduction

Indoor air quality (IAQ) is a critical factor in ensuring the health, comfort, and productivity of building occupants. As people spend approximately 80% of their time indoors (Klepeis et al., 2001), maintaining healthy indoor environments is increasingly important in both residential and commercial settings. Achieving this requires accurate modeling of indoor airflow. By capturing the full spatiotemporal dynamics of air movement, such models can identify poorly ventilated areas ("dead zones") that simplified models based on ordinary differential equations (ODEs) often overlook (He & Gonzalez, 2016; Hosseinloo et al., 2023). Furthermore, accurate airflow modeling supports smarter ventilation strategies that reduce unnecessary energy use while maintaining indoor air quality (Bian et al., 2024).

Computational Fluid Dynamics (CFD) has long served as a benchmark for high-fidelity indoor airflow simulation (Stamou & Katsiris, 2006), accurately capturing airflow behavior under varying room geometry and boundary conditions. These simulations solve partial differential equations (PDEs) using classical methods such as finite element, finite volume, or spectral techniques (Michoski et al., 2020). However, this accuracy comes at a high computational cost. Even with modern solvers, simulating airflow dynamics at room scale can take hours or even days on high-performance computing infrastructure. As a result, CFD remains primarily confined to offline applications—such as HVAC system design, layout validation, and safety assessment (Bulińska & Buliński, 2017; Bianco et al., 2023; Li et al., 2024)—and is rarely used for real-time or adaptive building control.

To bridge this gap, we leverage high-fidelity CFD simulations to build a data-driven surrogate model capable of approximating indoor airflow and CO_2 dynamics with orders-of-magnitude speed-up. Specifically, we simulate a realistic classroom environment with detailed ceiling-mounted HVAC configurations and varying occupancy scenarios. This allows us to generate a rich CFD dataset capturing spatiotemporal flow and CO_2 patterns under diverse boundary control conditions.

Building on this dataset, we propose the **Ensemble Neural Operator Transformer (ENOT)**—a novel operator learning framework that learns the mapping from boundary control and occupancy inputs to full-field CO_2 distributions over time. ENOT combines the expressive power of neural operators with ensemble modeling for improved generalization and uncertainty robustness. Its fast inference speed makes it suitable for real-time integration into downstream tasks such as model-based control and online optimization. Our key contributions are:

(1) We develop and release a high-fidelity CFD dataset capturing indoor air dynamics under varied HVAC control and occupancy settings in a real-world classroom.

(2) We propose ENOT that learns mappings from bound-

[1]Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, USA. Correspondence to: Yuexin Bian <yubian@ucsd.edu>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

ary control and occupancy inputs to spatiotemporal CO_2 distributions, achieving fast predictions with high accuracy.

(3) We highlight the broader applicability of our framework to several future directions, including data-driven model-based HVAC control, machine learning–accelerated building simulation, algorithm benchmarking, and a wide range of CFD-driven optimization tasks such as design and placement optimization.

2. CFD-Based Modeling and Dataset

2.1. Governing Equations for CO_2 Dynamics

Indoor CO_2 dynamics are governed by complex coupled physical processes, including airflow governed by the Navier–Stokes equations, CO_2 transport described by the advection–diffusion equation, and source terms arising from human occupancy. To accurately model these effects, we develop a high-fidelity CFD simulation in ANSYS Fluent (Manual, 2009) to capture the spatiotemporal evolution of CO_2 concentration. As shown in Fig 1, our simulation is based on a real-world classroom measuring approximately $19 \times 13 \times 3.5$ meters, equipped with a ceiling-mounted ventilation system. The system consists of 2 outlet vents and 18 inlet vents, organized into six zones. Each zone supports independent control of supply airflow rate and airflow angle, enabling energy-efficient operation while maintaining indoor air quality.

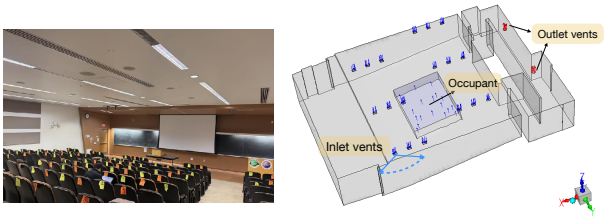

Figure 1. The picture of the studied classroom and its corresponding geometrical model used for CFD simulation.

Let $t \in \mathbb{R}^+$ denote time, and let $C(\mathbf{x}, t)$ represent the CO_2 concentration at spatial location $\mathbf{x} \in \Omega \subset \mathbb{R}^3$. We define the control action at time t as $m(t) = \begin{bmatrix} m^r(t) & m^a(t) \end{bmatrix} \in \mathbb{R}^{12}$, where $m^r(t), m^a(t) \in \mathbb{R}^6$ are the airflow rates and supply angles for the six zones, respectively. The occupancy level is denoted by $n_p(t) \in \mathbb{N}$, where $\mathbb{N} = \{0, 1, 2, ..\}$. The CO_2 distribution in an indoor environment follows the advection-diffusion equation (Bulińska & Buliński, 2017):

$$\frac{\partial C(\boldsymbol{x},t)}{\partial t} + \boldsymbol{u}(\boldsymbol{x},t) \cdot \nabla C(\boldsymbol{x},t) = D_{\text{eff}} \nabla^2 C(\boldsymbol{x},t) + S(\boldsymbol{x},t), \tag{1}$$

where $\boldsymbol{u}(\boldsymbol{x}, t)$ is the airflow velocity field obtained from CFD simulations, by solving the incompressible Navier-Stokes equations. The supply vent boundary conditions are determined by the building control input $m(t)$, allowing it to influence $\boldsymbol{u}(\boldsymbol{x}, t)$ across the domain. D_{eff} is the CO_2 diffusion coefficient, and $S(\mathbf{x}, t)$ models CO_2 emissions from occupants, with an exhalation rate of 6 L/min per person (He et al., 2022). For the occupancy boundary, we define $S(\boldsymbol{x}, t) = 6 \cdot n_p(t)$ L/min.

2.2. Dataset Generation

To capture a diverse set of physically realistic CO_2 distribution patterns, we simulate a wide range of control and occupancy conditions by sampling inlet airflow rates, airflow angles, and occupancy levels from uniform distributions:

$$m_i^r \sim U[\underline{m^r}, \overline{m^r}],\ m_i^a \sim U[45^\circ, 135^\circ],\ i \in [1, \ldots, 6],$$
$$n_p \sim U[10, 80],$$

where $U(\cdot, \cdot)$ denotes a uniform distribution, m_i^r is the airflow rate for the i-th group of vents, bounded between $\underline{m^r}$ = 0.324m/s (10% of maximum) and $\overline{m^r}$=3.24m/s; m_i^a is the airflow angle of the i-th group of vents, spanning 45° to 135°; and n_p is the number of occupants in the classroom.

All simulations solve the incompressible Navier–Stokes equations for airflow dynamics, coupled with species transport equations to model the distribution of CO_2, O_2, H_2O, and N_2. Turbulence effects are captured using the k–ω SST (Shear Stress Transport) model (Abuhegazy et al., 2020). The governing equations are discretized using the finite volume method, with second-order schemes applied to both momentum and species transport. Boundary conditions include velocity inlets and pressure outlets for the ventilation system, localized mass-flow inlets to represent human CO_2 exhalation, and no-slip conditions on all walls to realistically capture surface friction and flow resistance. The simulations consist of two phases: **(1) Steady-state initialization:** 10 simulations with randomized parameters were run to generate equilibrium airflow and CO_2 fields, serving as initial conditions for the next phase. **(2) Transient simulation:** 300 simulations were conducted by randomly selecting a steady-state solution and sampling control parameters m_i^r, m_i^a, and occupancy n_p from uniform distributions. CO_2 evolution was then simulated and recorded every 30 seconds over $\overline{T} = 60$ time steps (30 minutes).

To support further research, we release the simulation dataset, as detailed in Table 1. CO_2 concentrations were monitored at two heights: the HVAC level (2.9m) near the inlets to assess ventilation quality, and the people surface (1.6m) at typical standing head height (ASHRAE, 1992).

3. Operator Learning for Indoor Airflow Dynamics Modeling

Simulating a single trajectory with CFD can take hours to days, depending on mesh resolution, room/building size, and boundary control conditions. This motivates the use of

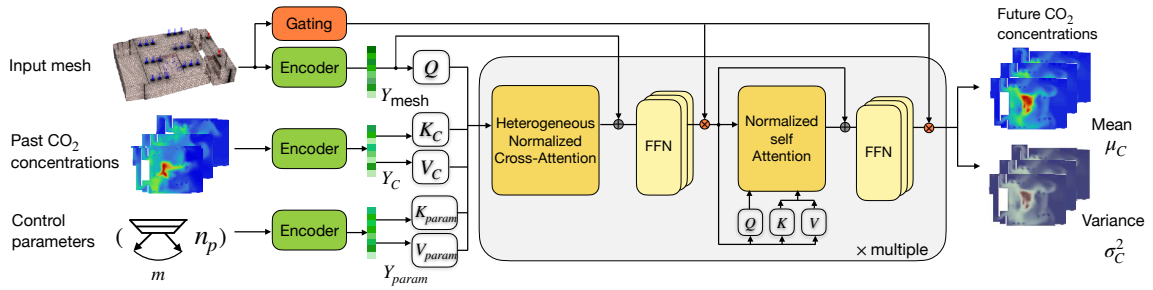

Figure 2. The neural operator transformer $\mathcal{G}_\theta$. We train an ensemble of Nensemble models and obtain the final prediction using Eq. (4).

Table 1. Data fields.

Field	Description
HVAC surface	`ndarray` (N_{hvac}, 3), spatial coordinates of grid points on HVAC surfaces
CO_2-HVAC	`ndarray` (N_{hvac},$\overline{T}$), CO_2 concentration time series at HVAC surface
People surface	`ndarray` (N_{people}, 3), spatial coordinates of grid points on people surfaces
CO_2-People	`ndarray` (N_{people},$\overline{T}$), CO_2 concentration time series at people surface
Steady case	`int`, Identifier from 0 to 9 for the initial steady-state used in the simulation
n_p	`int`, number of occupants
m_i^r	`float`, airflow rate (m^3/s) for i-th group of vents
m_i^a	`float`, angle ($^\circ$) for i-th group of vents

a surrogate model that can mimic the underlying physics with much lower computational burden, while retaining accuracy for control optimization. However, most neural network architectures for PDEs learn mappings between finite-dimensional Euclidean spaces (Sun et al., 2020), limiting predictions to a fixed set of spatial locations (Gao et al., 2024). In contrast, neural operators (Lu et al., 2021; Li et al., 2021; 2024; Hao et al., 2023) learn mappings between infinite-dimensional function spaces, producing *discretization-invariant* representations of PDE solutions that can be evaluated at arbitrary locations within the domain, enabling modeling of indoor air quality across the entire room. Neural operators have proven effectiveness for learning PDE solutions, particularly in fluid dynamics (Li et al., 2021; 2024).

In this work, we propose ENOT, which extends the General Neural Operator Transformer (GNOT) (Hao et al., 2023) with ensemble learning to improve accuracy for learning indoor dynamics.

3.1. Ensemble Neural Operator Transformer

We consider PDEs defined over a spatial domain $\Omega \subset \mathbb{R}^3$. A neural operator $\mathcal{G}_\theta$ is trained to learn a mapping $\mathcal{G} : \mathcal{A} \to \mathcal{H}$, where $\mathcal{A}$ is the input function space and $\mathcal{H}$ is the output function space. The input space $\mathcal{A}$ includes initial conditions, boundary conditions, source terms, and system parameters and output space $\mathcal{H}$ includes the PDE solutions. The neural operator $\mathcal{G}_\theta$ approximates the PDE solution operator $\mathcal{G}$,

$$\mathcal{G} : (C(\boldsymbol{x}, t-H:t), m, n_p) \mapsto C(\boldsymbol{x}, t:t+T) \qquad (2)$$

where $C(\boldsymbol{x}, t-H:t)$ is the historical CO_2 fields over period $(t-H, t]$ and $C(\boldsymbol{x}, t:t+T)$ is the predicted future CO_2 fields over the future interval $(t, t+T]$. Our forecasting approach incorporates historical CO_2 concentrations to account for temporal dependencies inherent to the system's physics (e.g., diffusion and advection dynamics). We further assume that the control m and occupancy n_p remain fixed over $(t, t+T]$, as building controls and occupancy typically remain constant over short intervals despite rapid transient airflow dynamics and CO_2 transport. In our approach, we assume that occupancy is known in advance over the control horizon. While this may not hold in all indoor environments, it is a reasonable assumption in many practical settings such as classrooms, conference rooms, or hospital wards, where occupancy follows predefined schedules or can be inferred from booking systems. This assumption allows the controller to anticipate CO_2 generation and plan ventilation accordingly. Extending our framework to account for stochastic or uncertain occupancy is a promising direction for future work.

The architecture of the ENOT is shown in Figure 2. To accommodate these heterogeneous inputs, a general encoder, highlighted in green in Figure 2, is employed to transform them into the feature embedding $Y \in \mathbb{R}^{N \times n_e}$, where N denotes an arbitrary number of input elements and n_e is the embedding dimension. The model employs simple multilayer perceptrons (MLPs), denoted as f_{w1}, f_{w2}, f_{w3}, to map each type of input to its corresponding embedding: Mesh points $\boldsymbol{x}_i$ are mapped to query embeddings: $Y_{\text{mesh}} = (f_{w1}(\boldsymbol{x}_i))_{1 \le i \le N_x} \in \mathbb{R}^{N_x \times n_e}$, where N_x is the number of mesh points, and n_e is the embedding dimension. At time t, historical CO2 concentrations $c_i = C(\boldsymbol{x}_i, t-H:t)$ at each location $\boldsymbol{x}_i$ are jointly encoded by another MLP, yielding $Y_C = (f_{w2}(\boldsymbol{x}_i, c_i))_{1 \le i \le N_x} \in \mathbb{R}^{N_x \times n_e}$. Control parameters $[m, n_p] \in \mathbb{R}^{13}$ are embedded into $Y_{\text{param}} = f_{w3}([m, n_p]) \in \mathbb{R}^{1 \times n_e}$. After encoding, inputs are pro-

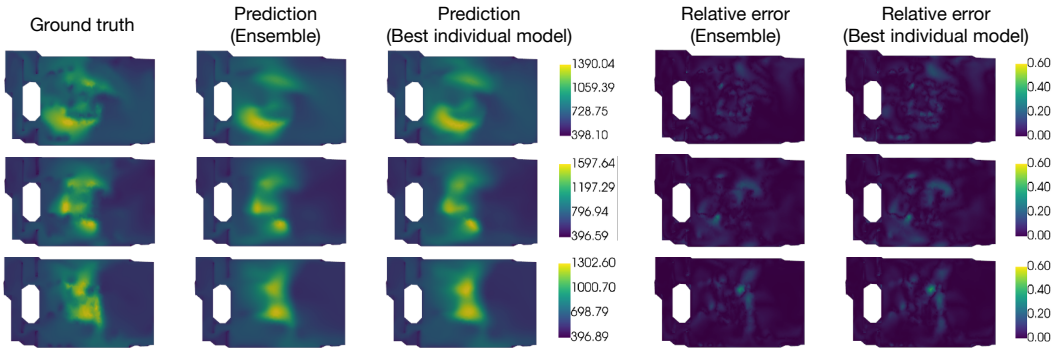

Figure 3. CO_2 predictions (ppm) and relative errors for three test cases. The ensemble model shows improved spatial accuracy and lower error.

cessed by a heterogeneous normalized cross-attention layer, followed by a self-attention layer. To effectively capture spatial heterogeneity, GNOT incorporates a geometric gating mechanism that leverages the query point coordinates to compute a weighted combination of expert feed-forward networks (FFNs). The model stacks N such attention blocks to generate predictions.

While GNOT is originally designed to predict only the mean of the prediction, we enhance its robustness by introducing an *ensemble learning* extension. Specifically, we modify the GNOT output to predict a probability distribution for future CO_2 concentrations, characterized by the mean (μ_C) and variance (σ_C^2), where $[\mu_C]_i^t, [\sigma_C^2]_i^t$ represents the mean and variance at location i and time step t. The model is trained using the Negative Log-Likelihood (NLL) loss:

$$\mathcal{L} = \frac{1}{N_x T} \sum_{i=1}^{N_x} \sum_{k=1}^{T} \left(\frac{\log(2\pi[\sigma_C^2]_i^{t+k})}{2} + \frac{(C_i^{t+k} - [\mu_C]_i^{t+k})^2}{2[\sigma_C^2]_i^{t+k}} \right) \tag{3}$$

where $C_i^{t+k} := C(\boldsymbol{x}_i, t+k) \in \mathbb{R}$ is the groundtruth CO_2 value at $\boldsymbol{x}_i$ and time $t+k$ from the CFD simulation. By minimizing the NLL loss, the model learns to jointly optimize the mean and variance, effectively mitigating overfitting (Lakshminarayanan et al., 2017). We further improve performance by training an ensemble of neural operator transformers. Each model $\mathcal{G}_{\theta_n}$ produces its own mean and variance prediction, and the ensemble output is computed as a uniformly weighted mixture, where N_{ensemble} is the number of models, and $\widehat{C} := \widehat{C}(\boldsymbol{x}, t\!:\!t\!+\!T)$ is the final prediction of the spatial-temporal CO_2 concentration,

$$\widehat{C} = \mu_C^{\text{ensemble}} = \frac{1}{N_{\text{ensemble}}} \sum_{n=1}^{N_{\text{ensemble}}} \mu_C^{(n)}. \tag{4}$$

3.2. Results

We train our ensemble GNOT to predict CO_2 concentrations on the people surface, using $N_x = 7462$ query points, with $H = 12$ input steps and $T = 6$ prediction steps. The dataset is split into 80% training and 20% testing. We train an ensemble of $N_{\text{ensemble}} = 5$ models and average their outputs. Each model is trained for 200 epochs using the AdamW optimizer with a cyclical learning rate schedule. Figure 3 compares the ground truth CO_2 concentration with predictions from the ensemble neural operator model and the best-performing individual model at the final prediction step, along with their corresponding relative error maps. These results highlight the neural operator framework's ability to model complex spatial CO_2 distributions under varying control conditions. Among the models, the ensemble GNOT model achieves the lowest test error at 10.90%, outperforming the five individual models, whose errors range from 11.82% to 13.01%.

In terms of efficiency, our model significantly outperforms traditional CFD simulations. We summarize the runtime associated with each stage of our operator learning pipeline in Table 2. All simulations were run on a Linux system with 16 CPU cores and an NVIDIA GeForce RTX 2080 Ti GPU. The CFD requires 1,253.7 seconds to compute $T = 6$ transient flow steps, while our model takes just 0.005 seconds. This represents a remarkable speed-up of approximately 250,000 times compared to the CFD approach. Such a dramatic reduction in computational time opens the potential for real-time control of complex ventilation systems, with the accurate airflow PDE models.

Table 2. Runtime of the operator learning pipeline.

Stage	**Time Cost**
CFD data generation	3.5 CPU hours per simulation, totaling 1100 CPU hours
Neural operator training	16 GPU hours
CFD simulation (6 steps)	1253.7 seconds
Inference (6 steps)	<0.005 seconds

4. Application and Outlooks

4.1. Data-Driven HVAC Control

Our work demonstrates the potential of operator learning to model indoor airflow using high-fidelity CFD data. Traditional HVAC control methods, such as model predictive control (MPC), often rely on simplified or black-box models that fail to capture complex spatiotemporal airflow patterns. In contrast, our learned model can be integrated into control frameworks to optimize ventilation, reduce energy consumption, and maintain indoor air quality. We have demonstrated the effectiveness of our approach (Bian & Shi, 2025), where we compare ENOT against state-of-the-art modeling methods, including reduced-order model (ROM)-based approaches and machine learning (ML)-based average prediction models. Our results show that ENOT achieves superior performance in capturing the control-response relationship. Furthermore, this improved modeling accuracy translates to better downstream control performance, highlighting the benefits of operator learning in building control tasks.

4.2. Machine Learning Accelerated Building Simulation

Unlike traditional CFD-based simulations that are computationally expensive and rely on hand-crafted features, our framework enables fast, data-driven simulations to support advanced building management. By comparing model predictions with real-time sensor data, it can quickly detect anomalies such as sensor or actuator failures, allowing for timely diagnosis and mitigation of system disruptions.

4.3. Algorithm Performance Benchmark

Our CFD-based simulation dataset serves as a benchmark for evaluating data-driven airflow modeling methods. By capturing spatiotemporal dynamics under varying control conditions, it enables fair comparisons between operator learning and other surrogate modeling approaches. In future work, we plan to expand the dataset to include additional rooms and support closed-loop evaluation.

4.4. Generalization to CFD-Based Optimization

Beyond HVAC control, our operator learning framework can be extended to a broad class of CFD-driven optimization problems, such as placement optimization and design optimization. These tasks typically require running a large number of computationally expensive CFD simulations to explore the design space. In contrast, our approach leverages a limited set of high-quality CFD data to learn a surrogate model that captures the underlying physics. This learned model can then be used to efficiently evaluate new design configurations, enabling gradient-based or black-box optimization without the need to exhaustively simulate every possibility. By replacing costly CFD evaluations with fast and accurate operator predictions, our framework opens the door to scalable and data-efficient design workflows in complex fluid systems.

5. Conclusion

In this work, we release a high-fidelity CFD dataset and propose ENOT, a neural operator-based framework for modeling indoor airflow and CO_2 concentration. Our method achieves high predictive accuracy while offering orders-of-magnitude speed-up compared to traditional CFD solvers, making it suitable for real-time applications and scalable optimization.

Beyond accurate modeling, we demonstrate how ENOT enables data-driven, model-based control by integrating with control algorithms to optimize ventilation and maintain indoor air quality with greater energy efficiency. We also highlight its potential as a benchmark environment for evaluating learning-based control algorithms in building systems.

Looking ahead, we envision several promising directions for extending this work. First, we aim to generalize our framework to multi-zone and larger-scale buildings, which pose additional modeling and control challenges. Second, we plan to integrate ENOT into real-world building management systems (BMS), enabling robust and adaptive HVAC control under practical deployment constraints. Finally, we see opportunities to apply our framework to a wider range of CFD-based optimization problems—such as design and placement optimization—by leveraging limited high-quality simulation data to accelerate the design loop without resorting to exhaustive CFD evaluations.

6. Impact Statement

This paper aims to advance the field of machine learning for building system modeling and control. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Abuhegazy, M., Talaat, K., Anderoglu, O., and Poroseva, S. V. Numerical investigation of aerosol transport in a classroom with relevance to covid-19. *Physics of Fluids*, 32(10), 2020.

ASHRAE. ASHRAE Standard 55–thermal environmental conditions for human occupancy. Technical report, ASHRAE Inc., Atlanta, GA, 1992.

Bian, Y. and Shi, Y. Data-driven operator learning for energy-efficient building control. *arXiv preprint arXiv:2504.21243*, 2025.

Bian, Y., Fu, X., Gupta, R. K., and Shi, Y. Ventilation and temperature control for energy-efficient and healthy buildings: A differentiable pde approach. *Applied Energy*, 372:123477, 2024.

Bianco, N., Fragnito, A., Iasiello, M., and Mauro, G. M. A cfd multi-objective optimization framework to design a wall-type heat recovery and ventilation unit with phase change material. *Applied Energy*, 347:121368, 2023.

Bulińska, A. and Buliński, Z. A cfd analysis of different human breathing models and its influence on spatial distribution of indoor air parameters. *Computer Assisted Methods in Engineering and Science*, 22(3):213–227, 2017.

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

Hao, Z., Wang, Z., Su, H., Ying, C., Dong, Y., Liu, S., Cheng, Z., Song, J., and Zhu, J. Gnot: A general neural operator transformer for operator learning. In *International Conference on Machine Learning*, pp. 12556–12569. PMLR, 2023.

He, R. and Gonzalez, H. Zoned hvac control via pde-constrained optimization. In *2016 American Control Conference (ACC)*, pp. 587–592. IEEE, 2016.

He, Y., Chu, Y., Zang, H., Zhao, J., and Song, Y. Experimental and cfd study of ventilation performance enhanced by roof window and mechanical ventilation system with different design strategies. *Building and Environment*, 224: 109566, 2022.

Hosseinloo, A. H., Nabi, S., Hosoi, A., and Dahleh, M. A. Data-driven control of covid-19 in buildings: a reinforcement-learning approach. *IEEE Transactions on Automation Science and Engineering*, 2023.

Klepeis, N. E., Nelson, W. C., Ott, W. R., Robinson, J. P., Tsang, A. M., Switzer, P., Behar, J. V., Hern, S. C., and Engelmann, W. H. 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.

Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

Li, Z., Kovachki, N. B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., Anandkumar, A., et al. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

Li, Z., Kovachki, N., Choy, C., Li, B., Kossaifi, J., Otta, S., Nabian, M. A., Stadler, M., Hundt, C., Azizzadenesheli, K., et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36, 2024.

Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G. E. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.

Manual, U. Ansys fluent 12.0. *Theory Guide*, 67, 2009.

Michoski, C., Milosavljević, M., Oliver, T., and Hatch, D. R. Solving differential equations using deep neural networks. *Neurocomputing*, 399:193–212, 2020.

Stamou, A. and Katsiris, I. Verification of a cfd model for indoor airflow and heat transfer. *Building and Environment*, 41(9):1171–1181, 2006.

Sun, L., Gao, H., Pan, S., and Wang, J.-X. Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Computer Methods in Applied Mechanics and Engineering*, 361:112732, 2020.