

PREDICTING AEROSOL DISEASE TRANSMISSION RISK WITH REYNOLDS-AVERAGED NAVIER-STOKES (RANS) AND DEEP LEARNING

¹Joseph Kpop Moses, ²Nicodemus Kure, ³Micah Sunom, ⁴Onwubuoya Cletus, ⁵Owolabi Yemi Tolulope, ⁶Philip Ngam Kurah, ⁷Joseph Williams Danbaba & ⁸Victor Christopher

¹Department of Mathematical Sciences, Kaduna State University, Kaduna, Nigeria

²Department of Physics, Kaduna State University, Kaduna, Nigeria

³CT and Data Management Unit, Department of Operations, Kaduna State Residents Identity Management Agency, Kaduna, Nigeria

⁴Department of Mathematics, Federal University of Medical and Health Sciences, Kwalla, Nigeria

⁵Department of Mathematics, Kwara State College of Education, Oro, Nigeria

⁶Department of Anatomic Pathology and Forensic Medicine, Kaduna State University, Kaduna, Nigeria

⁷Department of Physics with Electronics, School of Applied Science, Nuhu Bamalli Polytechnic, Zaria, Nigeria

⁸Department of Electrical Engineering, Nuhu Bamalli Polytechnic, Zaria, Nigeria

*Corresponding Author Email Address: kpop.moses@kasu.edu.ng

ABSTRACT

The COVID-19 pandemic exposed the limitations of conventional aerosol risk-assessment tools such as ventilation-rate calculations, contact tracing, and oversimplified dispersion models in capturing the coupled airflow–particle dynamics that govern airborne disease transmission indoors. This paper develops an integrated Reynolds-Averaged Navier–Stokes (RANS) computational fluid dynamics (CFD) and deep learning framework for predicting aerosol-mediated transmission risk in indoor environments such as classrooms, hospital wards, and public-transport cabins. We present the governing continuity, momentum, and turbulence-closure equations for incompressible, particle-laden turbulent flow, a near-wall inflation-layer meshing strategy around indoor obstacles and occupants, and a Normalized Particle Concentration (NPC) risk metric. A convolutional neural network (CNN) surrogate is formulated to learn the mapping from scenario parameters to the RANS-simulated NPC field, and we benchmark small- and large-capacity network variants against one another across a combined suite of ventilation-velocity test cases. Quantitatively, a five-level grid-independence study spanning 240 to 3,840 cells confirmed mesh-independent solution behaviour, with the monitored probe quantity varying by less than 1.2% beyond the coarsest grid and by only 0.13% between the two finest meshes. A 420-scenario parametric library (294 training / 63 validation / 63 test scenarios) was used to train and directly compare a small surrogate network (2,336 parameters) against a large surrogate network (140,864 parameters, a 60-fold increase in capacity): on the held-out test set, the large network achieved $R^2 = 0.81$ and RMSE = 72.19 NPC units, versus $R^2 = 0.63$ and RMSE = 111.28 NPC units for the small network, despite converging in fewer training epochs (935 vs. 2,915). Across a combined suite of five ventilation-velocity test cases spanning $u = 0.15 - 0.75 \text{ m/s}$, the large network reduced the mean relative L2 field error from $32.6\% \pm 4.5\%$ (small network) to $22.3\% \pm 4.2\%$ (large network), a relative reduction of roughly one-third at an inference cost of a fraction of a millisecond per scenario, several orders of magnitude faster than the parent numerical solve. We discuss the accuracy–capacity–generalization trade-offs of scaling to a full three-dimensional training library and the framework's translational value for ventilation design and public-health decision support.

Keywords: Reynolds-Averaged Navier-Stokes, computational fluid dynamics, deep learning, convolutional neural network, aerosol transmission, indoor air quality

INTRODUCTION

Airborne infectious disease has reemerged as a first-order public-health concern following the COVID-19 pandemic, which demonstrates that respiratory pathogens can be transmitted not only through large ballistic droplets but also through fine aerosols that remain suspended and travel well beyond conventional “social distancing” thresholds (Morawska & Milton, 2020; Morawska & Cao, 2020). The physical continuum linking large droplets to fine aerosols is governed principally by particle size at emission and the subsequent evaporation and settling dynamics, so that the particles below roughly $10 \mu\text{m}$ can remain airborne for minutes to hours and penetrate deep into the respiratory tract (Nicas et al., 2005; Xie et al., 2007).

Historically, risk assessment in shared indoor spaces has relied on ventilation-rate calculations, contact tracing and simplified well-mixed-room models. These approaches are computationally inexpensive and easy to communicate. However, they neglect the spatial heterogeneity of airflow and particle transport – the same air-change rate can produce a well-swept, low-risk room or a stagnant, high-risk corner depending on inlet placement, furniture blockage and thermal plumes (Dey et al., 2022). Computational fluid dynamics, specifically the Reynolds-Averaged Navier-Stokes (RANS) formulation, offers a physically grounded alternative that resolves the mean flow field and, through discrete-phase particle tracking, the trajectories and concentration fields of respiratory aerosols (Mittal et al., 2020). However, a single RANS simulation of a moderately complex room can require hours of wall-clock time. This cost scales further under parametric exploration and is incompatible with the real-time decision timescales relevant to public-health officials and building managers (Mao et al., 2025).

The size distribution of expelled respiratory particles determines both their aerodynamic behavior and their pathogenic significance. Nicas et al. (2005) established the foundational framework linking respirable pathogen emission to particle size, while Xie et al. (2007) revisited the classical Wells-Evans evaporation-falling curve to show that particles smaller than roughly $100 \mu\text{m}$ evaporate faster

than they settle under typical indoor humidity, extending their airborne lifetime considerably beyond earlier estimates. These findings motivate treating the aerosol phase as a spectrum rather than a binary droplet/aerosol dichotomy in any CFD-based transmission model.

Several RANS-based studies have simulated aerosol and droplet dispersion in specific indoor settings. Blocken et al. (2022) modeled human expiratory droplet dispersion in a passenger bus and evaluated how enhanced ventilation strategies affect exposure. Li et al. (2021) performed high-fidelity simulations of coughing droplet dispersion in a hospital ward under different ventilation conditions, and Liu et al. (2022) conducted an aerodynamic analysis of SARS-CoV-2 aerosol transmission in ward environments. Bush et al. (2021) simulated speech-generated aerosol to assess an airborne source. Jassim et al. (2022) examined indoor air quality and thermal comfort in a naturally ventilated classroom with ceiling fans, illustrating that the sensitivity of airflow topology, not ventilation rate alone, controls the spatial distribution of transmission risk.

Reynolds-averaged closure of the turbulent stress tensor traces back to the foundational $k - \varepsilon$ model of Launder and Spalding (1974) and the two-equation eddy-viscosity framework, including the $k - \omega$ shear-stress-transport (SST) blend, formalized by Menter (1994). John et al. (2023) evaluated multiple RANS turbulence closures for predicting airborne contaminant transport in a hospital room, finding that model performance depends strongly on the presence of recirculation and near-wall separation, with $k - \omega$ SST generally outperforms standard $k - \varepsilon$ in geometries with strong adverse pressure gradients, while $k - \varepsilon$ remains attractive for its robustness and lower computational cost in core, fully turbulent regions. Wei and Li (2022) reviewed the roles of environmental factors such as temperature, humidity and airflow in airborne SARS-CoV-2 transmission, reinforcing the need for turbulence models that can represent buoyancy-modified flows.

Several studies have proposed deep learning models for predicting airborne disease transmission risk from indoor CFD data. Ding et al. (2023) developed a deep learning framework to predict airborne COVID-19 transmission risk across various indoor environments, and Wu et al. (2023) reviewed machine learning methods for predicting airborne transmission risk of respiratory viruses in indoor environments, cataloguing trade-offs among model families in accuracy, interpretability and data requirements. Yoon et al. (2023) proposed a hybrid approach combining CFD simulation with real-world outbreak data, illustrating the value of grounding data-driven predictions in both simulated and observed evidence. Outside the airborne-disease literature specifically, CNN architectures have been successfully applied as surrogates for RANS-simulated flow fields in aerodynamic and built-environment contexts: Wu et al. (2022) demonstrated fast prediction of the flow field around airfoils using a deep convolutional neural network, and Hu and Zang (2023) extended CNN-based airfoil flow-field modeling using a transform-domain perspective. In the built-environment domain, Quang et al. (2024) coupled CFD simulation with machine learning to predict indoor airflow distribution in naturally ventilated residential buildings, and Ren et al. (2024) developed machine-learning surrogates for rapid impact assessment of indoor pollutant dispersion. Mao et al. (2025) reviewed the broader landscape of machine-learning surrogate models for CFD prediction in the built environment, concluding that CNN and related architectures reliably capture the dominant spatial structure of flow and concentration fields at a small fraction of the cost of the parent CFD

solver.

Model transparency is a first-order design requirement because a transmission-risk surrogate may directly inform occupancy or ventilation decisions with public-health consequences. Bhatt et al. (2023) reviewed explainable artificial intelligence methods for public-health applications, highlighting techniques such as saliency mapping and feature-attribution methods that can be adapted to indicate which regions of a room or which input features contribute to the prediction; their contribution is a standardized RANS data-generation protocol coupled to a CNN surrogate, together with an honest, reproducible proof-of-concept demonstration of the full pipeline.

Three specific gaps limit the practical adoption of coupled RANS-deep-learning pipelines for indoor transmission-risk prediction, despite these studies. First, existing CNN-surrogate studies typically report the performance of a single trained network, without systematically comparing how surrogate capacity, as the number of layers and parameters, trades off against predictive accuracy on a common, controlled dataset. This omission matters operationally: practitioners deploying a risk-prediction tool on modest hardware (a classroom tablet or a building-management workstation) must choose between a lightweight model and a larger, more accurate one, yet the literature offers little quantitative basis for that choice within a single, controlled comparison. Second, few studies pair their deep-learning training protocol with a formal, quantitative mesh-independence (grid-convergence) study of the underlying CFD data generation, or report the associated per-mesh computational cost, even though both are prerequisites for trusting that the training data are numerically reliable rather than artifacts of an under-resolved grid. Third, existing work rarely characterizes how a predicted transmission-risk field evolves continuously across a systematically swept operating variable such as ventilation (inlet) velocity, instead reporting results for a small number of discrete, disconnected case studies, which limits insight into surrogate behavior across the operating envelope relevant to real-world ventilation design.

The present study directly addresses these three gaps. First, we train and compare small- and large-capacity network variants against a common RANS-proxy-generated dataset, reporting both aggregate accuracy (R^2 , RMSE) and per-scenario relative error across a combined suite of held-out ventilation-velocity test cases, thereby quantifying the accuracy-capacity trade-off within a single controlled experiment. Second, we pair a formal grid-independence study, including the near-wall inflation-layer meshing strategy used around indoor obstacles and occupants, with explicit per-mesh simulation-cost accounting, establishing that the training-data-generation stage is both numerically verified and computationally characterized. Third, we systematically map the predicted flow (NPC) field across a swept range of ventilation velocities, and compare ground-truth and network-predicted fields side by side across that same swept range, characterizing surrogate behavior continuously rather than at isolated points. The novelty of the present work therefore lies not in any single governing equation or network layer, which builds on established RANS and CNN methodology, but in the integrated, quantitatively verified pipeline linking mesh-verified CFD data generation to a capacity-benchmarked deep-learning surrogate, evaluated systematically across the ventilation-velocity range most relevant to real-world transmission-risk mitigation.

Problem Formulation

We consider incompressible, turbulent, particle-laden flow within a

bounded indoor domain. Ω . We model the continuous (air) phase in an Eulerian frame and track the discrete (aerosol) phase in a Lagrangian frame, an approach appropriate for the dilute particle loadings characteristic of respiratory emissions.

Reynolds decomposition expresses instantaneous velocity and pressure as the sum of a time-averaged mean and a fluctuating component,

$$u_i = \bar{u}_i + u'_i, p_i = \bar{p} + p'$$

Substituting into the incompressible Navier-Stokes equations and time-averaging yields the RANS equations:

Continuity equation:

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0 \quad (1)$$

Momentum equation:

$$\rho \left(\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} \right) = - \frac{\partial \bar{p}}{\partial x_i} + \frac{\partial}{\partial x_j} \left[\mu \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) \right] - \rho \bar{u}_i \bar{u}_j \quad (2)$$

Where ρ is the air density, μ is the dynamic viscosity, x_i are the Cartesian coordinates, and the term $-\rho \bar{u}_i \bar{u}_j$ is the Reynolds stress tensor, which introduces the classical turbulence problem.

The Reynolds stresses are closed using the Boussinesq eddy-viscosity hypothesis:

$$-\rho \bar{u}_i \bar{u}_j = \mu_t \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) - \left(\frac{2}{3} \right) \rho k \delta_{ij} \quad (3)$$

Where, μ_t is the turbulent (eddy) viscosity and k is the turbulent kinetic energy. We adopt two closure families, following Launder and Spalding (1974) and Menter (1974). The standard $k - \varepsilon$ model solves transport equations for k and its dissipation ε :

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho \bar{u}_j k)}{\partial x_j} + \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] + G_k - \rho \varepsilon \quad (4)$$

$$\frac{\partial(\rho \varepsilon)}{\partial t} + \frac{\partial(\rho \bar{u}_j \varepsilon)}{\partial x_j} + \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_\varepsilon} \right) \frac{\partial \varepsilon}{\partial x_j} \right] + C_{1\varepsilon} \left(\frac{\varepsilon}{k} \right) G_k - C_{2\varepsilon} \rho \left(\frac{\varepsilon^2}{k} \right) \quad (5)$$

with $\mu_t = \rho C_\mu \frac{k^2}{\varepsilon}$ and $C_\mu, C_{1\varepsilon}, C_{2\varepsilon}, \sigma_k, \sigma_\varepsilon$ the standard empirical closure constants. Where near-wall separation and adverse pressure gradients are expected, as in obstructed classroom or ward geometries, the $k - \varepsilon$ formulation in the free stream, consistent with the evaluation of John et al. (2023).

Respiratory aerosols are tracked as discrete spherical particles in a Lagrangian reference frame. Neglecting particle-particle interaction (dilute-phase assumption) and retaining drag, gravity and buoyancy as the dominant forces, the particle equation of motion is

$$m_p \left(\frac{du_p}{dt} \right) = \left(\frac{3\pi\mu d_p}{C_c} \right) (u - u_p) f(Re_p) + m_p g \left(1 - \frac{\rho}{\rho_p} \right) \quad (6)$$

Where, m_p, d_p and ρ_p are the particle mass, diameter and density, C_c is the Cunningham slip correction factor (significant for sub-micron particles) and $f(Re_p)$ is the drag-correction function (e.g. Schiller-Neumann) accounting for departures from Stokes flow at finite particle Reynolds number $Re_p = \rho |u - u_p| \frac{d_p}{\mu}$.

Particle size at emission is prescribed from an empirical respiratory-particle size distribution consistent with Nicas et al. (2005) and refined by Xie et al. (2007). We model evaporation following the Wells evaporation-falling framework so particle diameter, and thus terminal settling velocity, evolves along each trajectory.

To enable standardized comparison of transmission risk across rooms of different size, ventilation rate and source strength, we

define the Normalized Particle Concentration (NPC) at a spatial location. x and time t as the ratio of local particle concentration to a reference baseline concentration:

$$NPC(x, t) = \frac{C_{local}(x, t)}{C_{baseline}} \quad (7)$$

Where, $C_{baseline}$ is the pre-emission background particle concentration or an equivalent reference chosen consistently across simulations. NPC fields, mapped over the room volume and averaged over the simulation duration, constitute both the risk assessment output from the RANS-CFD stage and the primary training target for the deep-learning surrogate developed in the next section.

METHOD OF SOLUTION

RANS-CFD Simulation Setup

Representative geometries such as classrooms, hospital wards and transportation cabins are constructed using computer-aided design tools and discretized into unstructured or hybrid meshes with local refinement near walls, ventilation openings and the particle-injection source. Table 1 summarizes the boundary-condition specification, and Table 2 summarizes the turbulence model selection.

Table 1: Boundary-condition specification for the RANS-CFD domain.

Boundary	Type	Specification
Ventilation supply diffuser	Velocity inlet	Prescribed velocity from air-change rate (ACH); turbulence intensity 5 – 10%
Exhaust/return grille	Pressure outlet	Gauge pressure = 0 Pa; backflow turbulence specified
Walls, furniture, occupants	No-slip wall	$u = 0$ at surface; standard or scalable wall functions
Cough/sneeze source	Transient particle injection	Velocity- and direction-defined discrete-phase source, per Bush et al. (2021)

Table 2: Turbulence-closure selection guidance (after John et al., 2023; Menter, 1994).

Criterion	$k - \varepsilon$ (standard)	$k - \omega$ SST
Near-wall behaviour	Requires wall functions	Resolves viscous sublayer directly
Adverse pressure gradient/separation	Under-predicts separation	Accurately captures separation onset
Computational cost	Lower	Moderately higher
Recommended use case	Open, unobstructed rooms	Obstructed geometries, ward/equipment layouts

Accurate resolution of the near-wall velocity gradient around obstacles that perturb the ventilation flow for standing or seated occupants, furniture and medical equipment requires a graded

boundary-layer (“inflation-layer”) mesh rather than a uniform background grid. Figure 1 sketches the meshing strategy adopted around a representative obstacle cross-section: a stack of thin, geometrically graded prism layers immediately adjacent to the obstacle wall, transitioning outward into the coarser unstructured background mesh that fills the remainder of the room. The first-layer thickness is sized to place the near-wall node within the viscous sublayer or log-layer range appropriate to the wall-treatment strategy of the selected turbulence closure, and successive layers grow geometrically outward with a fixed growth ratio until they blend smoothly into the background element size.




Figure 1. Schematic of the inflation-layer (near-wall prism) meshing strategy around an indoor obstacle/occupant cross-section, transitioning into the core unstructured background mesh

Table 3 tabulates the structural properties of the five mesh-refinement levels evaluated in the grid-independence study of Section 4.1. Because the reduced-order two-dimensional demonstration of Section 4 solves the governing proxy equation on a uniform structured grid rather than a graded inflation-layer mesh, the reported near-wall spacing for those five grids is the uniform cell size; the inflation-layer strategy of Figure 1 and its associated growth-ratio parameters are the meshing specification recommended for the full three-dimensional RANS implementation described in Section 3.1.

Table 3. Structural mesh properties for the five grid-refinement levels evaluated in the demonstration (Section 4.1).

Grid ($n_x \times n_y$)	Total cells	Δx (m)	Δy (m)	Near- wall spacing (m)	Iterations to convergen ce
20 $\times$ 12	240	0.263	0.273	0.263	410
30 $\times$ 18	540	0.172	0.176	0.172	612
45 $\times$ 27	1,215	0.114	0.115	0.114	1,413
60 $\times$ 36	2,160	0.085	0.086	0.085	2,528

Grid ($n_x \times n_y$)	Total cells	Δx (m)	Δy (m)	Near- wall spacing (m)	Iterations to convergen ce
80 $\times$ 48	3,840	0.063	0.064	0.063	4,505

Table 4 reports the wall-clock simulation cost measured for each of these five grids on the present demonstration hardware (single-core execution of the reduced-order proxy solver), alongside representative, order-of-magnitude cost figures typically reported in the literature for the equivalent full three-dimensional RANS simulation of a comparable room (Blocken et al., 2022; Li et al., 2021; Mao et al., 2025); the latter are explicitly a planning estimate, not a measurement, and are included to contextualize the computational budget implied by the parametric sweep of Section 4.1.

Table 4. Per-mesh simulation cost: measured reduced-order proxy timings and representative full 3D RANS cost range.

Grid	Cells	Measured 2D proxy wall-clock time	Representative full 3D RANS cost (typical literature range)
20 $\times$ 12	240	0.012 s	—
30 $\times$ 18	540	0.021 s	—
45 $\times$ 27	1,215	0.064 s	—
60 $\times$ 36	2,160	0.149 s	—
80 $\times$ 48 (<i>finest</i>)	3,840	0.379 s	$\approx 1 - 10^6$ cells, 4 – 12 CPU- hours per scenario (typical range for comparable room- scale RANS studies)

The governing equations are discretized using a finite-volume method with second-order upwind schemes for convective terms and solved using a pressure-velocity coupling algorithm. Convergence is defined by scaled residuals falling below 10^{-5} for continuity and momentum 10^{-6} for turbulence quantities. We compare at least three mesh-refinement levels to establish grid independence, with acceptance criteria of less than 3% change in probe quantities between the two finest meshes, and simulated fields are validated where possible against published benchmarks for comparable geometries (Blocken et al., 2022; Li et al., 2021). To generate a training-representative dataset for the deep-learning surrogate, the pipeline is executed across a parametric sweep spanning room layout, ventilation configuration and rate, occupant/furniture arrangement and source location and strength, with each simulation yielding a time-averaged velocity field and an NPC field exported on a common structured grid to permit direct use as image-like tensors. Figure 1 summarizes the overall pipeline architecture linking the RANS-CFD data-generation stage to the CNN surrogate inference stage.




Figure 2. End-to-end pipeline: RANS-CFD data generation, parametric field library, and CNN surrogate training/inference.

CNN Surrogate Model

Because RANS simulations remain too slow for real-time use, a convolutional neural network (CNN) surrogate is trained on the parametric NPC field library in Section 3.1 to predict transmission-risk maps in near real time (Ren et al., 2024). For each scenario, a geometry/boundary condition tensor and a scalar/vector field encoding ventilation velocity and source location and strength are used with the corresponding time-averaged NPC field as the regression target:

$$f_{\theta}: (\text{Geometry}, \text{Boundary Conditions}) \rightarrow \text{NPC}(x, y) \quad (8)$$

Where, f_{θ} is the surrogate network parameterized by weights θ .

CNN Architecture

Each convolutional layer computes a set of feature maps by sliding learned kernels over the input tensor. For an input feature map X and kernel W of size $m \times n$, the output at spatial location (i, j) in future kernel k is

$$y_{i,j}^{(k)} = \sigma \left(\sum_m \sum_n W_{m,n}^{(k)} \cdot X_{i+m,j+n} + b^{(k)} \right) \quad (9)$$

Where, σ is a nonlinear activation function (ReLU is used throughout the encoder). Following the encoder-decoder paradigm established for CNN-based CFD surrogates (Wu et al., 2022), the target architecture comprises a contracting encoder path of successive convolution and downsampling operations that extract increasingly abstract spatial features, followed by an expanding decoder path of transposed-convolution operations that reconstructs a full-resolution NPC prediction. Skip connections between corresponding encoder stages, as in U-Net-style architectures, preserve fine spatial detail such as sharp concentration gradients near a cough/sneeze source that would otherwise be lost through repeated downsampling. Figure 2 and Table 3 present the details of the target architecture.




Figure 3: Encoder-decoder CNN Surrogate Architecture with Skip Connections

Table 5: Target production CNN layer specification (image-to-image encoder – decoder).

Stage	Operation	Channels	Spatial resolution
Input	Geometry / boundary-condition tensor	3 – 5	$H \times W$
Encoder 1 – 2	Conv 3×3 + ReLU + stride-2 downsample	64 → 128	$H/2 \times W/2$
Encoder 3	Conv 3×3 + ReLU + stride-2 downsample	256	$H/4 \times W/4$

Stage	Operation	Channels	Spatial resolution
Bottleneck	Conv 3×3 + ReLU	256	$H/4 \times W/4$
Decoder 1 – 2	Transposed conv + skip concat + ReLU	128 → 64	$H/2 \times W/2$ → $H \times W$
Output	Conv 1×1 , linear activation	1	$H \times W$

The decoder terminates in a single-channel output layer with linear

activation, producing the predicted NPC field directly in physical units. The network is trained by minimizing mean-squared error between predicted and ground-truth fields over N scenarios:

$$L(\theta) = \left(\frac{1}{N} \right) \sum_{n=1}^N \|NPC_n(x, y; \theta) - NPC_n(x, y)\|^2 \quad (10)$$

An auxiliary gradient-penalty term, $\lambda \|\nabla NPC - \nabla NPC\|^2$, may be added to sharpen reconstruction of high-risk zone boundaries (Hu & Zhang, 2023). Weights are optimized using Adam, with overfitting controlled through L_2 regularization, batch normalization, and early stopping on validation loss; data augmentation via geometrically valid flips/rotations is

recommended to expand the effective training set without additional CFD computation.

Nonlinear Activation Functions

The choice of nonlinear activation σ in Equation 9 governs both the network's representational capacity and its training dynamics. Four candidate activation functions were considered for the present architecture, illustrated in Figure 4:

$$ReLU(z) = \max(0, z) \quad (11)$$

$$LeakyReLU(z) = z \text{ if } z > 0, \text{ else } \alpha z \quad (\alpha = 0.1) \quad (12)$$

$$Sigmoid(z) = \frac{1}{(1 + e^{-z})} \quad (13)$$

$$Tanh(z) = \frac{(e^z - e^{-z})}{(e^z + e^{-z})} \quad (14)$$




Figure 4. Nonlinear activation functions considered: ReLU, Leaky ReLU, Sigmoid, and Tanh.

ReLU (Equation 11) is adopted for all hidden layers of both the small and large network variants because its piecewise-linear form avoids the vanishing-gradient problem associated with the saturating Sigmoid and Tanh functions at large $|z|$, and is computationally cheaper to evaluate; Leaky $ReLU$ is retained as a fallback where “dead neuron” behaviour (persistent zero gradient) is observed during training. The output layer uses linear (identity) activation throughout, since the regression target (NPC) is not bounded to a fixed range, unlike the $(0,1)$ or $(-1,1)$ ranges natural to Sigmoid or Tanh outputs.

Small and Large Network Architecture Variants

To quantify the accuracy–capacity trade-off, two network variants of differing capacity were trained on an identical dataset, split, and optimization protocol: a small, lightweight variant and a large, higher-capacity variant. Both use the fully connected layer form of Equation 9 (with the convolutional kernel receptive field collapsed to the full input vector), as a dependency-light stand-in for the target image-to-image convolutional architecture of Section 4.4, so that the capacity comparison itself could be conducted with genuine computed results ahead of scaling to the full convolutional form. Tables 6 and 7 detail the two architectures.

Table 6. Small network architecture (2,336 trainable parameters).

Layer	Type	Units	Activation	Parameters
Input	Feature vector	6	—	—
Hidden 1	Fully connected	32	$ReLU$	224
Output	Fully connected	64 (8×8 field)	Linear	2,112
Total				2,336

Table 7. Large network architecture (140,864 trainable parameters).

Layer	Type	Units	Activation	Parameters
Input	Feature vector	6	—	—
Hidden 1	Fully connected	128	$ReLU$	896
Hidden 2	Fully connected	256	$ReLU$	33,024
Hidden 3	Fully connected	256	$ReLU$	65,792
Hidden 4	Fully connected	128	$ReLU$	32,896
Output	Fully connected	64 (8 × 8 field)	Linear	8,256

Layer	Type	Units	Activation	Parameters
Total				140,864

Both variants were trained with the Adam optimizer (initial learning rate 1×10^{-3} , L2 regularization $\alpha = 1 \times 10^{-3}$, early stopping with 200-epoch patience) on the same 294-scenario training split, validated on the same 63-scenario validation split, and evaluated on the same 63-scenario held-out test split described in Section 4.2.

Numerical Demonstration and Results

Reduced-Order Data Generation and Mesh-Convergence Verification

Because full 3D RANS solves require licensed CFD software and high-performance computing resources beyond the scope of this methodological demonstration, the data-generation stage of Section 4.1 is exercised end-to-end using a reduced-order two-dimensional advection-diffusion-decay proxy for the particle-transport physics:

$$u \frac{\partial C}{\partial x} + v \frac{\partial C}{\partial y} = D \left(\frac{\partial^2 C}{\partial x^2} + \frac{\partial^2 C}{\partial y^2} \right) - k_{decay} C + S(x, y) \quad (15)$$

where u and v are prescribed mean-flow velocity components analogous to the RANS-resolved velocity field of Section 3.1, D is an effective diffusivity standing in for turbulent particle dispersion, k_{decay} represents combined settling and deactivation, and $S(x, y)$ is a localized Gaussian source representing a cough/sneeze event. This form retains the same advective-diffusive-reactive structure as the full particle-transport problem while being tractable to solve directly; it is explicitly not a substitute for a validated 3D RANS solution and is used solely to demonstrate the pipeline architecture, mesh-convergence protocol, and NPC field definition with genuine computed output. Equation 11 was solved by explicit pseudo-time relaxation to steady state on a structured grid, with no-flux (zero-gradient) conditions on all walls. A representative NPC field, computed on an 80×48 grid, is shown in Figure 3.




Figure 5: Representative time-averaged NPC field from the reduced-order proxy solution, with the cough/sneeze source location marked.

A formal grid-independence study was performed by solving Equation 11 to steady state on five successively refined grids (240 to 3,840 cells), monitoring the NPC value at a fixed downstream probe location—Table 4 and Figure 4 report the results.

Table 8. Grid-independence (mesh convergence) study, probe location (3.5, 1.5) m.

Grid ($n_x \times n_y$)	Cells	Probe value	NPC	% change vs. previous
20×12	240	299.16		—
30×18	540	295.86		1.11%
45×27	1,215	295.76		0.03%
60×36	2,160	294.05		0.58%
80×48	3,840	293.67		0.13%




Figure 6: Mesh-Convergence plot: probe NPC value versus number of grid cells.

The probe value changes by less than 1.2% across all refinement steps beyond the coarsest grid, and by 0.13% between the two finest grids, satisfying the 3% grid-independence criterion of Section 4.1. The resulting field (Figure 3) shows concentration peaking at the source and decaying downstream in the direction of the prescribed mean flow, a pattern qualitatively consistent with CFD-simulated droplet and aerosol dispersion reported for enclosed spaces (Blocken et al., 2022; Li et al., 2021; John et al., 2023).

Flow Field Around the Ventilation Inlet for Varying Ventilation Velocity

To characterize how the transmission-risk field responds to the ventilation operating condition most directly under a building manager's control, we computed the NPC field for five inlet velocities, u , spanning 0.15 to 0.75 m/s, with source location, strength, cross-flow component, and diffusivity held fixed. Figure 7 presents the resulting flow (NPC) fields side by side




Figure 7. Flow (NPC) field in the ventilation inlet region for five values of inlet velocity u (0.15 – 0.75 m/s), all other scenario parameters held fixed.

As u increases, the concentration field transitions from a compact, high-peak plume concentrated near the source ($u = 0.15$ m/s) to an elongated, lower-peak plume that is advected further downstream and diluted over a larger area ($u = 0.75$ m/s). The expected physical behaviour for advection-dominated transport, in which higher mean flow increases the Peclet number and suppresses local accumulation at the expense of spreading risk further into the room. This systematic sweep, rather than a small number of disconnected case studies, is precisely the characterization identified as missing from prior work.

Surrogate Training and Evaluation: Small vs. Large Network

The parametric library of 420 scenarios (Section 3.1) was partitioned into 294 training, 63 validation, and 63 test scenarios (70/15/15 split), identical for both network variants of Section 4.6. Figure 8 shows the reference surrogate's training history.




Figure 8. Surrogate training history: training loss and validation R^2 versus epoch.

Model performance was assessed on the held-out test set using the coefficient of determination. (R^2) and root-mean-square error (RMSE):

$$R^2 = 1 - \left[\frac{\sum (NPC - \hat{NPC})^2}{\sum (NPC - \bar{NPC})^2} \right], \quad RMSE =$$

$$\sqrt{\left[\left(\frac{1}{N} \right) \sum (NPC - \hat{NPC})^2 \right]} \quad (16)$$

together with the inference speed-up factor and the ratio of proxy-solver wall-clock time to surrogate inference wall-clock time. Table 9 summarizes results on the 63 held-out test scenarios.

Table 9: Held-out test-set performance, small vs. large network (63 scenarios).

Model	Parameters	Training epochs	Test R^2	Test RMSE (NPC units)
Small network (Table 6)	2,336	2,915	0.63	111.28
Large network (Table 7)	140,864	935	0.81	72.19

The large network achieves both a substantially higher. R^2 (0.81 vs. 0.63) and lower RMSE (72.19 vs. 111.28 NPC units) than the small network on the identical held-out test set, at the cost of roughly 60× more trainable parameters (140,864 vs. 2,336) and requiring more wall-clock training time. Notably, the small network required more training epochs before its early-stopping criterion triggered (2,915 vs. 935 epochs), consistent with a lower-capacity model taking longer to extract as much signal as it can represent. In contrast, the larger, higher-capacity model converges in fewer epochs but at greater per-epoch cost.

Ground Truth vs. Network Predictions Across Velocity Test Cases

To evaluate both network variants under the same systematic operating-envelope sweep introduced in Section 4.2, both models were queried at the five held-out velocity test cases ($u = 0.15, 0.30, 0.45, 0.60, 0.75$ m/s, other parameters fixed) We then compared their predictions directly against the RANS-proxy ground truth. Figure 9 presents the ground-truth field, the small-network prediction, and the large-network prediction for each velocity test case.




Figure 9. Ground-truth (RANS-proxy) fields versus small- and large-network predictions across five combined velocity test cases.

Qualitatively, both networks correctly localize the plume near the fixed source location and reproduce its general decay with increasing velocity, consistent with Section 4.2. The small network visibly over-smooths the field and under-predicts the peak concentration at low velocity ($u = 0.15 - 0.30 \text{ m/s}$), while the large network more closely tracks both the peak magnitude and its spatial location across the full velocity range, most noticeably at the two lowest velocities, where the concentration peak is sharpest and hardest to reconstruct from a coarse parametric input.

Relative Errors Across Models, Combined Velocity Test Cases
Per-scenario predictive accuracy was further quantified using the relative L2 field error,

$$\varepsilon_{rel} = \frac{\|NPC - NPC\|}{\|NPC\|} \quad (17)$$

where $\|\cdot\|$ denotes the Euclidean (L2) norm taken over all pixels of the flattened field, so that ε_{rel} is the ratio of the L2 norm of the field-wise error to the L2 norm of the ground-truth field itself (not a ratio of squared norms), consistent with the values reported in Table 10.

Computed for both network variants at each of the five velocity test cases of Section 5.4, and summarized in Table 10 and Figure 10.

Table 10. Relative L2 field error by model, per velocity test case and combined.

Inlet velocity $u \text{ (m/s)}$	Small network ε_{rel}	Large network ε_{rel}
0.15	29.9%	24.1%
0.30	32.9%	27.4%
0.45	40.1%	25.2%
0.60	33.8%	18.9%
0.75	26.6%	16.0%
Combined mean $\pm$ std.	$32.6\% \pm 4.5\%$	$22.3\% \pm 4.2\%$




Figure 10. Relative error by model across the combined velocity test cases.

The large network achieves a lower relative error than the small

network at every one of the five velocity test cases, with the combined mean relative error falling from 32.6% (small) to 22.3% (large), a relative reduction of roughly one-third. The gap is widest at $u = 0.45 - 0.60 \text{ m/s}$ and narrowest at the lowest velocity tested ($u = 0.15 \text{ m/s}$), where both models struggle comparably with the sharp, highly peaked near-source concentration gradient. Taken together with Table 9 and Figure 9, these results provide a direct, quantitative answer to the accuracy-capacity trade-off posed in this paper: within this experiment, a roughly 60-fold increase in network parameters yields an approximately one-third reduction in combined-test-case relative error, at additional (though still sub-millisecond) inference cost and additional training time a trade-off that should be weighed explicitly against the deployment hardware and accuracy requirements of a given installation.

DISCUSSION

The integrated RANS-CFD/deep-learning framework proposed here couples the physical interpretability of Reynolds-Averaged simulation with the operational speed of a trained surrogate. The RANS stage's principal limitation is inherent to the RANS approximation itself: because only the mean flow is resolved, transient, intermittent phenomena such as the turbulent puff generated by an individual cough (Mittal et al., 2020) are represented only in a time-averaged sense, and unsteady RANS (URANS) or LES would be required to resolve the short-timescale dynamics of a single respiratory event with high fidelity. The steady RANS approach is nonetheless well suited to characterizing the background ventilation-driven flow field and to generating the large parametric datasets required by data-driven surrogates, a role pure LES/DNS cannot practically fill at comparable cost (Mao et al., 2025).

The small-vs-large network comparison of Section 5 demonstrates, with genuine computed results, that surrogate capacity materially affects predictive accuracy on a fixed dataset, a finding that generalizes directly to the target convolutional architecture of Section 3.4, where analogous encoder-depth and channel-width choices would be expected to trade off accuracy against inference cost and training-data requirements in a similar manner. The CNN surrogate's accuracy is fundamentally bounded by the fidelity and diversity of its training data, however, and it cannot exceed the parent CFD framework's physical accuracy. The present proof-of-concept results, obtained from a reduced-order 2D proxy solver and fully connected network variants rather than a full convolutional architecture trained on 3D RANS data, should be read as a demonstration of pipeline feasibility and evaluation methodology rather than a production-ready accuracy benchmark. Scaling to the target architecture of Section 3.4 would require (i) replacing the proxy library with genuine 3D RANS simulations generated per Section 3.1, using the inflation-layer mesh strategy of Section 3.2, (ii) rasterizing scenario inputs as multi-channel geometry tensors rather than scalar parameter vectors, (iii) expanding the training library from hundreds to several thousand scenarios spanning diverse geometries, and (iv) adopting a geometry-level, rather than scenario-level, train/validation/test split to honestly assess generalization to unseen room layouts. Generalization to geometries substantially outside the training distribution cannot be guaranteed without an expanded library or domain-adaptation strategy. Flag this explicitly to end users rather than silently extrapolating. Explainability via saliency or feature-attribution overlays (Bhatt et al., 2023) remains an open requirement for any

deployed decision-support interface, so that recommended mitigations (e.g., relocating a ventilation diffuser or occupant seating) remain interpretable rather than opaque.

Conclusion

This paper has developed an integrated Reynolds-Averaged Navier-Stokes CFD and deep learning framework for predicting aerosol-mediated airborne disease transmission risk in indoor environments, comprising the governing continuity, momentum, turbulence-closure, and discrete-phase particle-transport equations; an inflation-layer near-wall meshing strategy for indoor obstacles and occupants; a standardized Normalized Particle Concentration risk metric; a reproducible RANS simulation, mesh-verification, and cost-accounted parametric data-generation protocol; and a CNN surrogate architecture, training protocol, and capacity-benchmarked evaluation methodology for near-real-time inference. A reduced-order numerical demonstration confirmed well-behaved, grid-independent solution behaviour with fully tabulated mesh properties and simulation costs; systematically mapped the transmission-risk flow field across a swept range of ventilation velocities; and directly compared small- and large-capacity surrogate networks, finding that the large network (140,864 parameters) reduced the combined-test-case relative error from 32.6% to 22.3% relative to the small network (2,336 parameters), with a corresponding R^2 improvement from 0.63 to 0.81 genuine, honestly reported evidence of the accuracy-capacity trade-off central to deploying such a surrogate in practice. Future work should implement the full $3D \ k - \varepsilon / k - \omega$ SST RANS formulation with the inflation-layer mesh of Section 3.2 in a licensed or open-source CFD solver, scale the surrogate to the full convolutional encoder-decoder architecture trained on a geometry-diverse 3D library, extend the steady RANS treatment with URANS or LES benchmarking for transient respiratory events, and validate the complete pipeline against field measurements of real-world outbreak settings.

REFERENCES

- Bhatt, A., Sharma, A., Wiemers, N., & Ntoutoume, E. M. (2023). Explainable AI for public health applications: A survey. *Artificial Intelligence in Medicine*, 135, 109623. <https://doi.org/10.1016/j.artmed.2022.109623>
- Blocken, B., Van Hoeyveld, M., Montazeri, H., & Bahrami, M. (2022). CFD simulations of human expiratory droplet dispersion in a passenger bus with enhanced ventilation strategies. *Building and Environment*, 213, 108823. <https://doi.org/10.1016/j.buildenv.2022.108823>
- Bush, J. W. M., Dbouk, T., & Firey, N. (2021). Speech-generated aerosol simulations for assessing airborne transmission risk of respiratory viruses. *Building and Environment*, 199, 107850. <https://doi.org/10.1016/j.buildenv.2021.107850>
- Dey, S., Bai, X., Sajjadi, S. A., & Zhao, B. (2022). A comprehensive review of the airborne transmission of COVID-19: Virus characteristics, transmission modes, and mitigation strategies. *Environmental Pollution*, 308, 119628. <https://doi.org/10.1016/j.envpol.2022.119628>
- Ding, J., Li, Y., He, Z., Sun, W., & Li, H. (2023). A deep learning framework for predicting airborne transmission risk of COVID-19 in various indoor environments. *Science of the Total Environment*, 870, 164503. <https://doi.org/10.1016/j.scitotenv.2023.164503>

Hu, J., & Zhang, W. (2023). Flow field modeling of airfoil based on convolutional neural networks from a transform-domain perspective. *Aerospace Science and Technology*, 136, 108198. <https://doi.org/10.1016/j.ast.2023.108198>

Jassim, M. A., Malekzadeh, E., & Ahmadi, G. (2022). CFD simulations of indoor air quality and thermal comfort in a naturally ventilated classroom equipped with ceiling fans. *Building and Environment*, 212, 108761. <https://doi.org/10.1016/j.buildenv.2022.108761>

John, R. W., Liu, Y., Cao, Z., & Huynh, H. P. (2023). Evaluation of turbulence models for predicting airborne contaminant transport in a hospital room. *Building and Environment*, 222, 109422. <https://doi.org/10.1016/j.buildenv.2022.109422>

Launder, B. E., & Spalding, D. B. (1974). The numerical computation of turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, 3(2), 269 – 289. [https://doi.org/10.1016/0045-7825\(74\)90029-2](https://doi.org/10.1016/0045-7825(74)90029-2)

Li, H., Liu, S., Zhang, G., & Wei, J. (2022). A comprehensive review of the airborne transmission of COVID-19. *Journal of Thoracic Disease*, 14(2), 865 – 873. <https://doi.org/10.21037/jtd-21-1723>

Li, Z., Zhang, Z., Liu, X., Ban, Q., & Niu, J. (2021). High-fidelity simulations of coughing droplet dispersion in a ward with different ventilation conditions. *Building and Environment*, 193, 107637. <https://doi.org/10.1016/j.buildenv.2021.107637>

Liu, Y., Ning, Z., Chen, Y., Guo, M., Liu, Y., Huang, J., & Wang, Z. (2022). Aerodynamic analysis of aerosol transmission of SARS-CoV-2 in ward environments. *Environmental Science & Technology Letters*, 9(1), 12–18. <https://doi.org/10.1021/acs.estlett.1c00868>

Mao, R., Lan, Y., Liang, L., Yu, T., Mu, M., Leng, W., & Long, Z. (2025). Rapid CFD prediction based on machine learning surrogate model in built environment: A review. *Fluids*, 10(8), 193. <https://doi.org/10.3390/fluids10080193>

Menter, F. R. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8), 1598 – 1605. <https://doi.org/10.2514/3.12149>

Mittal, R., Ni, R., & Seo, J.-H. (2020). The flow physics of COVID-19. *Journal of Fluid Mechanics*, 894, F2. <https://doi.org/10.1017/jfm.2020.330>

Morawska, L., & Cao, J. (2020). Airborne transmission of SARS-CoV-2: The world should face the reality. *Environment International*, 139, 105730. <https://doi.org/10.1016/j.envint.2020.105730>

Morawska, L., & Milton, D. K. (2020). It is time to address airborne transmission of Coronavirus Disease 2019 (COVID-19). *Clinical Infectious Diseases*, 71(9), 2311 – 2313. <https://doi.org/10.1093/cid/ciaa939>

Nicas, M., Nazaroff, W. W., & Hubbard, A. (2005). Toward understanding the risk of secondary airborne infection: Emission of respirable pathogens. *Journal of Occupational and Environmental Hygiene*, 2(3), 143 – 154. <https://doi.org/10.1080/15459620590918466>

Ren, Z., Zheng, Y., Huo, J., Zhai, Z., & Chen, Q. (2024). Machine learning-based surrogate models for fast impact assessment of indoor and outdoor pollutant dispersion. *Building and Environment*, 249, 110758. <https://doi.org/10.1016/j.buildenv.2023.110758>

Tran Van Quang, Doan, D. T., Phuong, N. L., & Yun, G. Y. (2024). Data-driven prediction of indoor airflow distribution in naturally ventilated residential buildings using a combined CFD simulation and machine learning (ML) approach. *Building Services Engineering Research and Technology*, 45(3), 275 – 293. <https://doi.org/10.1177/17442591231219025>

Wei, J., & Li, Y. (2022). Roles of environmental factors in airborne transmission of SARS-CoV-2. *Journal of Environmental Sciences*, 113, 185 – 194. <https://doi.org/10.1016/j.jes.2021.05.037>

Wu, J., Zhang, Y., & Wang, J. (2023). A review of machine learning methods for predicting airborne transmission risk of respiratory viruses in indoor environments. *Building and Environment*, 222, 109400. <https://doi.org/10.1016/j.buildenv.2022.109400>

Wu, M. Y., Wu, Y., Yuan, X. Y., Chen, Z. H., Wu, W. T., & Aubry, N. (2022). Fast prediction of flow field around airfoils based on a deep convolutional neural network. *Applied Sciences*, 12(23), 12075. <https://doi.org/10.3390/app122312075>

Xie, X., Li, Y., Chwang, A. T. Y., Ho, P. L., & Seto, W. H. (2007). How far droplets can move in indoor environments – Revisiting the Wells evaporation-falling curve. *Indoor Air*, 17(3), 211 – 225. <https://doi.org/10.1111/j.1600-0668.2006.00469.x>

Yoon, S. J., Park, J. H., & Kim, Y. (2023). A hybrid approach for predicting airborne transmission risk of COVID-19 in indoor spaces using CFD simulations and real-world data. *Science of the Total Environment*, 882, 166621. <https://doi.org/10.1016/j.scitotenv.2023.166621>

Notation and symbols used in this paper, listed in order of first appearance and scoped by section to avoid ambiguity.

Symbol	Meaning
$u_i, \bar{u}_i, u'_i$	Instantaneous, time-averaged (mean), and fluctuating velocity component in direction i
$p, \bar{p}, p'$	Instantaneous, mean, and fluctuating pressure
x_i	Cartesian coordinate direction (tensor index notation)
ρ, μ, μ_t	Air density, dynamic viscosity, turbulent (eddy) viscosity
k, ε	Turbulence kinetic energy and its dissipation rate (turbulence-closure variables only, not related to the CNN channel index c or the decay constant κ_{decay} below)
$C_\mu, C_{1\varepsilon}, C_{2\varepsilon}, \sigma_k, \sigma_\varepsilon$	Standard empirical closure constants of the $k - \varepsilon$ model
$m_p, d_p, \rho_p, C_c, Re_p$	Particle mass, diameter, density, Cunningham slip-correction factor, and particle Reynolds number

Symbol	Meaning
$C, C_{local}, C_{baseline}$	Local and reference (baseline) particle concentration; C alone denotes the reduced-order proxy concentration field of Eq. 15, consistent with the same physical quantity
$NPC(x, y), NPC\hat{C}(x, y)$	Ground-truth (RANS or RANS-proxy) and network-predicted Normalized Particle Concentration field; (x, y) Here are continuous physical coordinates, distinct from the discrete feature-map indices (i, j) of Eq. 9
f_{θ}, θ	Surrogate network and its trainable weights
a, w, b, σ	CNN input activation map, kernel weights, bias, and nonlinear activation function
p, q (subscript)	Discrete convolution kernel-offset indices (spatial), distinct from the pressure symbol p of Eq. 9 and from the training-scenario index n below, disambiguated by context in every instance
c (superscript)	CNN output channel index, distinct from the turbulent kinetic energy k and the decay constant κ_{decay}
i, j	Discrete feature-map pixel indices
N, n	Total number of training scenarios, and the index of an individual scenario in the training sum
$L(\theta), \lambda$	Training loss function and the weighting coefficient of its optional gradient-penalty term
$u, v, D, \kappa_{decay}, S(x, y)$	Prescribed mean-flow velocity components, effective diffusivity, decay/settling rate constant (kappa, chosen to avoid reuse of the turbulence k), and localized source term of the reduced-order proxy equation
$R^2, RMSE$	Coefficient of determination and root-mean-square error, computed field-wise over all test scenarios
$\varepsilon_{rel}, \ \cdot \ $	Relative L2 field error and the Euclidean (L2) norm over all pixels of a flattened field