

Geometry-conditioned diffusion generation of Lagrangian particle trajectories in human airways

Ignacio Rodrigo Bartol

Georgia Institute of Technology

Martin Graffigna

Georgia Institute of Technology

Mauricio Tano

Idaho National Laboratory

Shaheen Dewji

`shaheen.dewji@gatech.edu`

Georgia Institute of Technology

Article

Keywords: Lagrangian particle tracking, human airways, diffusion probabilistic model, geometry conditioning, signed distance function, aerosol transport

Posted Date: August 4th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10503600/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Geometry-conditioned diffusion generation of Lagrangian particle trajectories in human airways

Ignacio R. Bartol^a, Martin Graffigna Palomba^a, Mauricio E. Tano^{b,*},
Shaheen A. Dewji^{a,*}

^a*Nuclear and Radiological Engineering and Medical Physics Program, George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, 770 State St NW, Atlanta, 30332-0405, GA, United States of America*

^b*Collaborative Computing Center (C3), Idaho National Laboratory, 955 MK Simpson Boulevard, Idaho Falls, 83415, ID, United States of America*

Abstract

Fast and subject-specific prediction of particle transport in the human airways is central to quantifying aerosol deposition, optimizing inhaled drug delivery, and supporting dose assessment in radiation research. While computational fluid and particle dynamics simulations provide high-fidelity trajectory data, their use in clinical or large-cohort settings remains limited by stringent geometric preprocessing (segmentation quality control, topology repair, surface smoothing, and meshing) and by the high computational cost of coupled airflow-particle tracking. Recent data-driven surrogates have accelerated aspects of airway flow prediction, but most approaches focus on Eulerian flow fields, operate in simplified or two-dimensional settings, and often struggle to generalize across complex three-dimensional anatomies. Especially for particle trajectories in a Lagrangian frame of reference, which are intrinsically multimodal and sensitive to local geometry.

In this work we propose a geometry-conditioned *guided denoising diffusion probabilistic model* that generates physically plausible particle trajectories directly in a Lagrangian representation. Airway anatomy is encoded by a variational autoencoder to provide a compact conditioning latent, and geometric feasibility during sampling is reinforced through signed distance function guidance that discourages samples leaving the airway lumen. Despite training on particle-track data from only nine simulated airway geometries, the proposed method produces trajectory ensembles that qualitatively match ground-truth branching patterns and recover the underlying velocity magnitude statistics with good agreement, while guidance reduces non-physical

*These authors jointly supervised this work

Email addresses: Co-Corresponding author: `mauricio.tanoretamales@inl.gov`
(Mauricio E. Tano), Co-Corresponding author: `shaheen.dewji@gatech.edu`
(Shaheen A. Dewji)

wall-penetrating outliers. Remaining limitations include occasional boundary violations in challenging cases, reduced diversity in some distal branches under strong guidance, and degraded behavior for out-of-distribution geometries not represented by the limited training set.

Keywords:

Lagrangian particle tracking, human airways, diffusion probabilistic model, geometry conditioning, signed distance function, aerosol transport

1. Introduction

Inhaled particle transport in the human respiratory tract spans a broad range of highly sensitive applications, from radioactive aerosol exposure assessment and respiratory toxicology to the design and optimization of inhaled therapeutics. For medical aerosols, therapeutic effectiveness depends on where and how much of the delivered dose deposits across the extrathoracic region, conducting airways, and peripheral lung, which is influenced by airway morphology, breathing waveform, and particle size distribution Labiris and Dolovich (2003); Longest and Holbrook (2012); Kundu and Panchagnula (2025); Kleinstreuer and Zhang (2010). As a result, subject-specific predictions are attractive for patient-tailored drug delivery and for identifying inter-subject variability in regional deposition Poorbahrami et al. (2019); Huang et al. (2021).

Beyond pharmacological aerosols, detailed particle tracking is central to inhalation dosimetry for radioactive aerosols found in occupational, environmental, and radiological accident scenarios ICRP (1994); Talaat et al. (2019); Bartol et al. (2024b). While regulatory and radiological protection frameworks often rely on population-average deposition models ICRP (1994), several studies have shown that high-resolution computational predictions can localize deposition hot spots and translate them into spatially resolved internal dose estimates that traditional models would have missed assuming uniform deposition across the airway lumen Talaat et al. (2019); Poorbahrami et al. (2019); Bartol et al. (2024b).

Nonetheless, modeling these aerosols is challenging. A multiphysics modeling approach is required; particle deposition is influenced by the interplay of inertial impaction, gravitational sedimentation, and Brownian diffusion, with the relative importance of each mechanism varying strongly with particle size, airway region, and local flow structures Darquenne (2014); Hofmann

(2011); Kleinstreuer and Zhang (2010). Consequently, high-fidelity simulations using Computational Fluid and Particle Dynamics (CFPD) remain too time-consuming for many clinical workflows that require rapid exploration of inhaler settings, breathing conditions, or patient anatomies.

State-of-the-art CFPD typically solves airflow using the Navier-Stokes equation in an Eulerian frame of reference, and tracks particles using a Lagrangian formulation, providing highly detailed, spatially resolved predictions of transport and particle deposition profiles Nowak et al. (2003); Longest and Holbrook (2012); Rostami (2009). Despite its accuracy and interpretability, constructing a subject- and scenario-specific CFPD pipeline is non-trivial. The workflow requires: segmentation of the airway tree from a medical imaging (computed tomography or magnetic resonance), 3D surface reconstruction and its repair (e.g., removal of disconnected branches, correction of holes and non-manifold regions), generation of a high-quality volumetric mesh, and careful specification of boundary conditions and turbulence modeling assumptions Longest and Holbrook (2012); Huang et al. (2021); Longest et al. (2019); Bartol et al. (2025). These steps are often stringent because small geometric defects (particularly near bronchi bifurcations) can produce ill-conditioned meshes or non-physical flow artifacts that propagate into particle deposition predictions Huang et al. (2021); Longest et al. (2019).

Even after preprocessing steps, computational wall time remains the primary bottleneck for computing a diverse dataset that covers multiple boundary conditions. High-fidelity predictions may require transient simulations, fine near-wall resolution, and large particle datasets to reduce variance and resolve rare events in the tail of deposition statistics Kleinstreuer and Zhang (2010); Longest et al. (2019); Huang et al. (2021). Consequently, while CFPD is well-suited for detailed analysis and device development, it is challenging to deploy at scale across large cohorts or embed in iterative clinical decision support.

Motivated by the cost of numerical simulation, machine learning surrogates for fluid dynamics have grown rapidly Brunton et al. (2020); Kochkov et al. (2021); Kutz and Brunton (2016). Early approaches using deep learning often encode geometry into a fixed grid and perform regression on Eulerian flow observables such as velocity or pressure fields Guo et al. (2016). More recent operator-learning approaches aim to learn mappings between function spaces, which are considered more interpretable and can make predictions across families of partial differential equations Lu et al. (2021); Li

et al. (2021); Huhn et al. (2023); Graffigna et al. (2026). Another approach used mesh-based graph neural networks, which provided a route to learning directly on unstructured discretizations Pfaff et al. (2021) and have the potential to overcome the limitation of other DL methods that worked only on structured grids. These advances meaningfully expand what is possible with data-driven models for computational physics Brunton et al. (2020); Kochkov et al. (2021), but they have been demonstrated only for Eulerian flow-field prediction, with limited translation to Lagrangian formulations.

Several gaps remain for subject-specific airway particle tracking. First, many learning-based surrogates are evaluated on simplified benchmarks or on problems with controlled geometry parameterizations, leaving open questions about robustness on heterogeneous, patient-derived three-dimensional airway topologies Brunton et al. (2020); Guo et al. (2016); Lu et al. (2021); Graffigna et al. (2026). Second, the dominant target in the literature is the Eulerian carrier-phase field (or aggregated deposition metrics), whereas the downstream quantity of interest in inhalation studies is often the particle-by-particle deposition profile and its uncertainty, including the exact distribution of particle tracks, residence times, and deposition outcomes Longest and Holbrook (2012); Huang et al. (2021); Graffigna et al. (2026). As explained, trajectory generation is intrinsically complex and produces high-variance distributions; even with fixed boundary conditions, the ensemble variability produced by turbulence, near-wall interactions, and stochastic forces (relevant for submicron particles) could be different for the same geometry, and the spatially accurate prediction is critical for risk assessment and dose estimation Darquenne (2014); Hofmann (2011). Finally, purely deterministic regressors as in Han et al. (2022) do not naturally provide a generative mechanism for sampling multiple plausible trajectory realizations conditioned on patient geometry and inhalation settings.

Diffusion probabilistic models address conditional generative modeling by learning a gradual denoising process that maps Gaussian noise to data Sohl-Dickstein et al. (2015); Ho et al. (2020); Nichol and Dhariwal (2021). In the denoising diffusion probabilistic model formulation, a forward Markov chain progressively perturbs data with noise, while a neural network learns the reverse-time denoising transitions Ho et al. (2020); Nichol and Dhariwal (2021). A common parameterization trains a model to predict the injected noise at diffusion step t as in Eq. 1, where $\mathbf{c}$ denotes conditioning information Ho et al. (2020); Nichol and Dhariwal (2021).

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \mathcal{L} = \mathbb{E}[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t, \mathbf{c})\|_2^2], \quad (1)$$

Compared to previous generative models, generative adversarial networks (GANs) Goodfellow et al. (2020), diffusion models are widely recognized for stable training (a supervised denoising objective rather than adversarial minimax optimization) and strong empirical sample quality; as it is well known, diffusion models have been shown to exceed state-of-the-art GANs on standard image synthesis benchmarks Dhariwal and Nichol (2021).

Notably, base DDPM models would not suffice to bridge the gap in subject specificity. Conditioning of DDPMs is of utmost importance in this setting, where the experiment requires samples consistent with known context, which for this work is the airway anatomy. Modern conditional diffusion models incorporate rich conditioning pathways, such as cross-attention modules and learned embeddings Rombach et al. (2022); Liu et al. (2023b). Moreover, guided diffusion enables steering of samples at inference time to enforce constraints without retraining the base model, particularly in medical or physical imaging settings, where hallucinations play a vital role Liu et al. (2023a). In image synthesis, guidance has been implemented using auxiliary gradients (classifier guidance) Dhariwal and Nichol (2021) and by combining conditional and unconditional scores, both of which are commonly used in practice and have been formalized in latent diffusion frameworks Rombach et al. (2022). These ideas lay the groundwork for subject-specific trajectory generation, where airway geometric plausibility can be viewed as a constraint that is otherwise difficult to guarantee from data alone (i.e., only particle trajectories).

Recent evidence also indicates that diffusion models are well-suited to learning Lagrangian particle trajectory statistics in complex dynamical systems: diffusion-based generative modeling has been shown to reproduce multi-scale statistics of three-dimensional turbulent single-particle trajectories, including heavy tails and intermittency effects Li et al. (2024a). This motivates exploring diffusion models as samplers for physically plausible trajectory ensembles beyond the turbulence setting, provided that appropriate conditioning and constraints are incorporated.

In this paper, we propose a geometry-conditioned guided DDPM for fast generation of airway particle trajectories. The overarching goal of this work is to replace repeated, computationally costly simulation-based trajectory gen-

eration with a learned conditional model that can produce plausible particle tracks for a given airway geometry in seconds to minutes of neural inference, while retaining key statistical-physical features of the underlying dynamics.

The proposed approach combines three ideas. First, we learn a compact latent representation of patient airway geometry with a *variational encoder* Kingma and Welling (2014) and use this latent vector as conditioning input to the denoising diffusion probabilistic model. Second, we represent airway shape using a signed distance function field – a representation that provides an explicit notion of walls and smooth geometric gradients Park et al. (2019); Mescheder et al. (2019). Third, during inference, we apply *signed distance function guidance* to discourage samples that violate geometric feasibility (e.g., trajectories leaving the airway lumen), analogously to guidance mechanisms used in conditional diffusion Dhariwal and Nichol (2021); Rombach et al. (2022).

Our contributions in this work are summarized as follows:

1. a diffusion-based generative model for trajectory ensembles conditioned on three-dimensional patient geometry,
2. a variational encoding strategy that enables conditioning on anatomies not seen during training,
3. an inference-time guidance mechanism based on signed distance function geometry that improves geometric plausibility,
4. an evaluation demonstrating that the model can recover velocity magnitude statistics and produce visually consistent trajectory bundles, while also revealing that generalization to unseen anatomies remains fundamentally limited by the diversity of available simulated geometries.

2. Methods

2.1. Study overview and scope

This study used high-fidelity simulation-derived particle tracks in a Lagrangian framework as the reference data for training and evaluation of a geometry-conditioned DDPM. The input anatomy for conditioning was represented as a 3D voxelized airway occupancy grid. The DDPM output was one fixed-length sequence of 3D particle positions per generated trajectory

192 (i.e. a sequence of particle tracks). The airway anatomy grid was encoded by
 193 a variational geometry encoder (VAE) and supplied to the temporal denoiser
 194 as global anatomical conditioning. A signed distance field (SDF), constructed
 195 from the same voxelized geometry but not used as a learned conditioning in-
 196 put, was queried during sampling to provide a soft geometric penalty for
 197 out-of-lumen positions at inference time.

198 The inlet flow, particle-size distribution, injection protocol, airway-simulation
 199 settings, and terminal handling rules were inherited from the simulation
 200 dataset rather than independently conditioned at inference. Therefore, the
 201 implementation reported in this manuscript approximates a conditional dis-
 202 tribution of position sequences given anatomy under a fixed data-generation
 203 regime and preprocessing regime represented in the training data. In its cur-
 204 rent version, it does not output carrier-flow variables, or a deposition/escape
 205 label. Comparisons with simulation-derived trajectories were made using the
 206 qualitative trajectory overlays and trajectory-derived kinematic metrics de-
 207 fined later in this section. In contrast to diffusion-based trajectory generators
 208 developed for unbounded homogeneous isotropic turbulence Li et al. (2024b),
 209 the present setting involves bounded, heterogeneous, subject-specific airway
 210 domains, which is a major challenge in real-world applications.

211 This separation between geometry representation, conditioning, and guid-
 212 ance is important for model interpretability. In one hand, the occupancy grid
 213 supplies a learned global geometric embedding, whereas the SDF is a directly
 214 queried geometric object to correct particle trajectories used during infer-
 215 ence. The reference particle-diameter distribution, injection parameters, and
 216 breathing conditions (waveform, intensity) defined the regime from which
 217 trajectories were sampled; in this study they were not separate conditioning
 218 variables for the DDPM as mentioned before. As a result, evaluation fo-
 219 cuses on whether generated position sequences follow the represented airway
 220 topology and reproduce a broad set of physical metrics focused on particle
 221 kinematics.

222 The remainder of this section describes the reference trajectory source,
 223 trajectory representation, geometric preprocessing, conditional diffusion model,
 224 SDF-guided sampling procedure, architecture and training protocol, metric
 225 definitions, and how we used dynamic mode decomposition (DMD) as base-
 226 line to compare different surrogate models. The carrier-phase and Lagrangian
 227 simulation information are included only facilitate the reading flow and pro-
 228 vide context for the fixed reference-data regime, the details on the simulation
 229 settings can be found in previous work from Bartol et al. (2025, 2024b).

230 2.2. Reference trajectory dataset

231 2.2.1. Large-cohort airway segmentation and representative downselection

232 Airway geometries were reconstructed from a large cohort of CT scans
 233 using the segmentation and quality-control workflow described previously
 234 Bartol et al. (2025). Briefly, low-quality scans, including scans with thick
 235 slices or incomplete thoracic coverage, were excluded before downstream
 236 geometry processing. The reconstruction pipeline produced 542 candidate
 237 lower-airway 3D reconstructed geometries as mesh surfaces. Automated geo-
 238 metric checks for non-physical branches, unrealistic dimensions, and related
 239 failures reduced this set to 412 anatomically valid lower-airway domains for
 240 population-level analysis via a downselection process explained below.

241 Each airway geometry i was summarized by a feature vector $\mathbf{z}_i$,

$$\mathbf{z}_i = [t_{d,i}, b_{a,i}, V_i, L_i, \dots]^\top, \quad (2)$$

242 where t_d is the volume-averaged tracheal diameter, b_a is the main-bronchi
 243 carina angle, V is the segmented airway volume, and L is the cumulative
 244 bronchial length. Feature vectors were standardized as

$$\tilde{\mathbf{z}}_i = \mathbf{D}^{-1}(\mathbf{z}_i - \boldsymbol{\mu}), \quad (3)$$

245 with $\boldsymbol{\mu}$ the cohort mean and $\mathbf{D}$ the diagonal matrix of feature standard
 246 deviations.

247 A Random Forest analysis in the prior downselection workflow identified
 248 t_d and b_a as informative descriptors for stratification. The reduced phenotype
 249 space was therefore defined as

$$\mathbf{y}_i = [\tilde{t}_{d,i}, \tilde{b}_{a,i}]^\top. \quad (4)$$

250 Phenotypic groups were then obtained using the clustering and representative-
 251 selection procedure explained in Bartol et al. (2024b). In the present manuscript,
 252 this process yielded nine male groups and nine female groups, containing the
 253 412 airway geometries. One representative geometry was selected from each
 254 group by choosing the nearest-neighbor exemplar to the reported cluster cen-
 255 troid,

$$i_k^* = \arg \min_{i \in \mathcal{C}_k} \|\mathbf{y}_i - \boldsymbol{\mu}_k\|_2, \quad \mathcal{G}_k^* \equiv \mathcal{G}_{i_k^*}. \quad (5)$$

256 resulting in a total of 18 representative geometries for the database.

257 As a summary, the downselection progression was therefore starting with
 258 542 candidate reconstructions, 412 anatomically valid lower-airway domains,
 259 18 representative geometries after sex-stratified downselection with available
 260 particle-track simulations. Only nine male geometries entered the present
 261 generative-model study. Consequently, all trajectory-generation evidence in
 262 this manuscript concerns the nine male simulated geometries rather than the
 263 full cohort of 412 anatomically valid geometries. The reason for not including
 264 the female geometries in this study, is that derived particle tracks from the
 265 female cohort is still under validation studies and were not yet published.

266 2.2.2. Carrier-phase airflow model

267 The reference carrier-phase airflow calculations used incompressible Reynolds-
 268 averaged Navier-Stokes (RANS) equations with the k - ω shear-stress-transport
 269 (SST) model Menter (1994) and the Langtry-Menter correlation-based tran-
 270 sition model (SST- γ - Re_θ , or SST-LM) Menter et al. (2006b); Langtry et al.
 271 (2004); Menter et al. (2006a). For a complete numerical please refer to the
 272 prior simulation work Bartol et al. (2025). The incompressible RANS equa-
 273 tions were written as

$$\nabla \cdot \bar{\mathbf{u}} = 0, \quad (6)$$

$$\rho \left(\frac{\partial \bar{\mathbf{u}}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} \right) = -\nabla \bar{p} + \nabla \cdot [(\mu + \mu_t) (\nabla \bar{\mathbf{u}} + \nabla \bar{\mathbf{u}}^\top)] + \mathbf{S}_p, \quad (7)$$

274 where $\bar{\mathbf{u}}$ is the mean velocity, $\bar{p}$ is the mean pressure, μ is the molecular
 275 viscosity, μ_t is the modeled eddy viscosity from the SST-LM closure, and $\mathbf{S}_p$
 276 is the particle back-coupling momentum source term described in Section 2.2.3.

277 A sinusoidal inlet-flow waveform representative of heavy exercise was pre-
 278 scribed at the tracheal inlet,

$$Q(t) = Q_{\max} \sin \left(\frac{2\pi t}{T} \right), \quad Q_{\max} = 100 \text{ L min}^{-1}, \quad T = 2 \text{ s}. \quad (8)$$

279 The inlet bulk velocity magnitude was then set from the instantaneous flow
 280 rate and inlet area A ,

$$\|\bar{\mathbf{u}}_{\text{in}}(t)\| = \frac{Q(t)}{A}. \quad (9)$$

281 Turbulence inlet quantities were specified from turbulence intensity I and
 282 a characteristic length scale L ,

$$k_{\text{in}} = \frac{3}{2} (I \|\bar{\mathbf{u}}_{\text{in}}\|)^2, \quad \omega_{\text{in}} = \frac{k_{\text{in}}^{1/2}}{C_\mu^{1/4} L}. \quad (10)$$

283 We set the initial turbulence intensity at $I = 4\%$, and L defined as the
 284 geometry inlet hydraulic diameter d_h . A value of 0.09 was set for the C_μ -
 285 related initialization constant.

286 2.2.3. Lagrangian particle tracking and two-way coupling

287 Particles were simulated using a Lagrangian frame with two-way coupling
 288 between the particles and carrier phase. For particle p with mass m_p , position
 289 $\mathbf{x}_p$, and velocity $\mathbf{v}_p$, the particle equation of motion was formulated as

$$\frac{d\mathbf{x}_p}{dt} = \mathbf{v}_p, \quad m_p \frac{d\mathbf{v}_p}{dt} = \mathbf{F}_p^{(\text{drag})} + \mathbf{F}_p^{(\text{grav/buoy})} + \mathbf{F}_p^{(\text{lift})} + \mathbf{F}_p^{(\text{vm})} + \mathbf{F}_p^{(\text{Brown})}. \quad (11)$$

290 The modeled forces included drag, gravity/buoyancy, shear lift, virtual
 291 mass, and Brownian forces in the submicron regime. The closures and coef-
 292 ficients follow the prior CFPD work Bartol et al. (2025).

293 Two-way coupling was represented by adding the reaction force from par-
 294 ticles to the Eulerian momentum equation as a volumetric source over each
 295 control volume Ω_c ,

$$\mathbf{S}_p(\Omega_c) = -\frac{1}{|\Omega_c|} \sum_{p \in \Omega_c} \mathbf{F}_{p \rightarrow f}, \quad (12)$$

296 where $\mathbf{F}_{p \rightarrow f}$ denotes the force exerted by a particle on the fluid, which is
 297 equal and opposite to the force applied to the particle.

298 Particle size distribution followed a log-normal sampling parameterized
 299 by an activity median aerodynamic diameter and a geometric standard de-
 300 viation, assuming spherical particles. For diameter d ,

$$p(d) = \frac{1}{d \ln(\sigma_g) \sqrt{2\pi}} \exp \left(-\frac{[\ln(d/d_g)]^2}{2 \ln^2(\sigma_g)} \right). \quad (13)$$

301 The values used for this work were $d_g = 0.42 \mu\text{m}$ and $\sigma_g = 3.5$.

302 Simulated particles were advanced in time until wall deposition, outlet es-
 303 cape, or the end of the simulation. Each exported reference track contained
 304 an event state in $\{\text{in-flight, deposited, escaped}\}$. In the DDPM implementa-
 305 tion, however, it was learned only to the three spatial position coordinates.
 306 Deposited and escaped reference trajectories were therefore represented in
 307 the training tensor through their terminal positions and the last-value-hold
 308 operation described in Section 2.3.1.

For each representative geometry $\mathcal{G}_k^*$, simulations were run under the fixed waveform described above, and particle tracks were exported as time-ordered sequences

$$\mathcal{T}^{(r)} = \{(t_\ell, \mathbf{x}_\ell, \mathbf{v}_\ell, d_p, \text{state}_\ell)\}_{\ell=1}^{L_r}, \quad (14)$$

where r indexes an exported trajectory and L_r varies with the termination or export time. The variables $\mathbf{v}_\ell$, d_p , and state_ℓ were retained in the reference export but were not part of the reported $C = 3$ diffusion-model output.

2.3. Trajectory representation and geometric preprocessing

2.3.1. Trajectory tensorization, time discretization, and variable-length handling

For training of the DDPM purposes, each exported trajectory $\mathcal{T}^{(r)}$ was resampled to a fixed length of $K = 1024$ physical-time indices,

$$t_k = k\Delta t, \quad k = 0, \dots, K-1, \quad K = 1024, \quad \Delta t = 0.00125 \text{ s}. \quad (15)$$

The learning state at time index k was the three-dimensional position,

$$\mathbf{s}_k \in \mathbb{R}^3, \quad \mathbf{x} = \{\mathbf{s}_k\}_{k=0}^{K-1} \in \mathbb{R}^{K \times 3}. \quad (16)$$

The $K \times 3$ notation was used for the mathematical description of particle trajectories. For one-dimensional convolutional implementation, batches were arranged in channel-first form $(B, 3, K)$. The final physical-time index is $(K-1)\Delta t = 1.27875 \text{ s}$, which for practical purposes is described throughout the manuscript as an approximately 1.28 s horizon. The justification to cut the simulation time from $t = 2 \text{ s}$ to $t = 1.28 \text{ s}$ is to avoid long-tail events since most of the particles were either deposited on the airway wall or escaped through an outlet at the selected cutoff time. Particle step size was selected relative to local cell size of approximately 1 mm and local particle velocity.

If a particle deposited on the airway wall or escaped through an outlet before t_{K-1} , the final valid position was extended by last-value hold,

$$\mathbf{x}_k = \mathbf{x}_{k^*} \quad \forall k \geq k^*, \quad (17)$$

where k^* is the last time index for which the trajectory position was updated. The approximately 1.28 s learning horizon was selected using an empirical histogram of last-updated times to reduce the amount of stationary terminal

padding relative to the full 2s waveform interval. Last-value hold creates zero displacement increments after termination and therefore contributed to near-zero finite-difference velocities and accelerations that later created an unrealistic representation in pooled kinematic distributions.

2.3.2. Coordinate frames and normalization

Each airway geometry was associated with voxelization metadata: an origin $\mathbf{o} \in \mathbb{R}^3$, voxel spacing $\mathbf{h} \in \mathbb{R}^3$, and grid dimensions $\mathbf{d} \in \mathbb{N}^3$. World coordinates were mapped elementwise to voxel coordinates by

$$\mathbf{x}^{(\text{vox})} = \frac{\mathbf{x}^{(\text{world})} - \mathbf{o}}{\mathbf{h}}, \quad (18)$$

and then normalized to the coordinate range used by the SDF query,

$$\mathbf{x}^{(\text{norm})} = 2 \left(\frac{\mathbf{x}^{(\text{vox})}}{\mathbf{d} - 1} \right) - 1. \quad (19)$$

The same normalized coordinate system was used for the position channels and for trilinear SDF sampling during guided diffusion. De-normalization to physical coordinates was applied for trajectory analysis.

2.3.3. Geometry voxelization and SDF construction

Each three-dimensional airway surface mesh was provided as an STL surface and voxelized to a binary lumen occupancy grid $\mathbf{B} \in \{0, 1\}^{D \times H \times W}$. The longest bounding-box dimension was fitted to a target maximum grid size, and zero-padding was applied if needed to produce the model input dimensions of $64 \times 64 \times 64$. The resulting 64^3 grid is an approximate representation of the detailed airway surface; it should not be interpreted as preserving every distal geometric feature of the original mesh. The binary occupancy grid was used for learned geometry conditioning.

A signed distance field $\mathbf{S} \in \mathbb{R}^{64 \times 64 \times 64}$ was computed from the voxelized representation and stored with the corresponding metadata $(\mathbf{o}, \mathbf{h}, \mathbf{d})$. The SDF was used only at inference time as a soft geometric guidance term. Throughout the manuscript, the adopted convention is

$$S(\mathbf{x}) < 0 \quad \text{inside the represented airway lumen}, \quad S(\mathbf{x}) > 0 \quad \text{outside the represented airway lumen}$$

Continuous SDF values were obtained by trilinear interpolation in the normalized coordinate frame used by Eq. (19).




Figure 1: **Binary occupancy and signed distance field for one airway geometry.** Top: slices of the binary voxelized lumen supplied to the variational geometry encoder. Bottom: slices of the signed distance field used during inference-time guidance. The adopted sign convention in this manuscript is negative inside the represented lumen and positive outside. Voxelization provides an approximate representation of the STL airway surface.

363 The two geometry representations therefore had distinct roles: **B** pro-
 364 vided the occupancy-based conditioning input to the denoiser, and **S** pro-
 365 vided an inference-time geometric penalty. Figure 1 illustrates the binary
 366 grid and corresponding SDF for one geometry.

367 2.4. Geometry-conditioned diffusion model

368 Let $\mathbf{x} \in \mathbb{R}^{K \times 3}$ denote a fixed-length three-dimensional particle-position
 369 track, and let $\mathcal{G}$ denote the airway geometrical field supplied through the
 370 occupancy grid. The model approximates a conditional distribution over
 371 position sequences, $p_{\theta}(\mathbf{x} \mid \mathcal{G})$, under the fixed inlet airflow waveform, particle
 372 size distribution, injection protocol, and preprocessing represented in the
 373 training data. The modeled object by the DDPM is the position trajectory
 374 over time. The training objective contains a denoising loss and a geometry-
 375 embedding KL regularization term.

376 Following diffusion probabilistic modeling Sohl-Dickstein et al. (2015);
 377 Ho et al. (2020); Song et al. (2021) and guided-diffusion implementations

378 used in related trajectory-generation work Li et al. (2024a); Dhariwal and
 379 Nichol (2021), a fixed Markov chain gradually corrupts a clean trajectory $\mathbf{x}_0$
 380 obtained from the CFPD simulation with Gaussian noise, the physical time
 381 index is denoted by k , whereas the diffusion step index is denoted by n . We
 382 can write the well known forward process as

$$q(\mathbf{x}_{1:N} | \mathbf{x}_0) = \prod_{n=1}^N q(\mathbf{x}_n | \mathbf{x}_{n-1}), \quad q(\mathbf{x}_n | \mathbf{x}_{n-1}) = \mathcal{N}(\sqrt{\alpha_n} \mathbf{x}_{n-1}, (1 - \alpha_n) \mathbf{I}). \quad (20)$$

383 Each noise corruption step retains a fraction $\sqrt{\alpha_n}$ of the previous signal
 384 and adds independent Gaussian noise with variance $\beta_n = 1 - \alpha_n$. The forward
 385 process is independent of the airway geometry; geometry enters through the
 386 learned reverse process and the inference-time guidance described below in
 387 this section.

388 While the derivation of the closed-form marginal is derived in the original
 389 work by Sohl-Dickstein et al. (2015), for any diffusion step can be written as

$$q(\mathbf{x}_n | \mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_n} \mathbf{x}_0, (1 - \bar{\alpha}_n) \mathbf{I}), \quad \Rightarrow \quad \mathbf{x}_n = \sqrt{\bar{\alpha}_n} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_n} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (21)$$

390 This closed-form allows training examples to be generated by sampling n
 391 directly rather than simulating every intermediate noising step, which would
 392 be much more inefficient.cf

393 The implementation used in this work was $N = 800$ diffusion steps and
 394 adopted the hyperbolic-tangent cumulative schedule `tanh6, 1` used in Li et al.
 395 (2024a),

$$\bar{\alpha}_n = \frac{-\tanh(7\frac{n}{N} - 6) + \tanh(1)}{-\tanh(-6) + \tanh(1)}, \quad n \in \{0, \dots, N\}, \quad (22)$$

396 which satisfies $\bar{\alpha}_0 = 1$ and $\bar{\alpha}_N \simeq 0$. The per-step coefficients were then
 397 obtained as

$$\alpha_n = \frac{\bar{\alpha}_n}{\bar{\alpha}_{n-1}}, \quad \beta_n = 1 - \alpha_n, \quad n = 1, \dots, N. \quad (23)$$

398 For the reverse denoising process, the generative model learns a reverse-
 399 time Markov chain that approximately inverts the forward corruption:

$$p_{\theta}(\mathbf{x}_{0:N-1} \mid \mathbf{x}_N, \mathcal{G}) = \prod_{n=N}^1 p_{\theta}(\mathbf{x}_{n-1} \mid \mathbf{x}_n, \mathcal{G}). \quad (24)$$

400 The idea of conditioning the neural network on $\mathcal{G}$ is to make the reverse
 401 process dependent on the anatomy, meaning that for a given random Gaus-
 402 sian noise initialization $\mathbf{x}_N$, the denoising process is now conditioned on the
 403 encoded airway representation so now $p_{\theta}(\mathbf{x}_{0:N-1} \mid \mathbf{x}_N) \rightarrow p_{\theta}(\mathbf{x}_{0:N-1} \mid \mathbf{x}_N, \mathcal{G})$.
 404 Each reverse transition was modeled as a Gaussian distribution,

$$p_{\theta}(\mathbf{x}_{n-1} \mid \mathbf{x}_n, \mathcal{G}) = \mathcal{N}(\boldsymbol{\mu}_{\theta}(\mathbf{x}_n, n, \mathcal{G}), \sigma_n^2 \mathbf{I}). \quad (25)$$

405 The reverse variance was fixed rather than learned as in previous litera-
 406 ture, where σ_n^2 is set from the forward schedule, and $\sigma_n^2 = \beta_n$ for most steps
 407 and the forward-posterior variance at the first reverse step.

408 The denoiser predicts the Gaussian noise added to a position-sequence
 409 tensor,

$$\hat{\boldsymbol{\epsilon}}_{\theta} = \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_n, n, \mathcal{G}), \quad (26)$$

410 where $\boldsymbol{\epsilon}_{\theta}$ is a one-dimensional temporal U-Net with the embedded condi-
 411 tioning. Given $\hat{\boldsymbol{\epsilon}}_{\theta}$, the new position estimate is

$$\hat{\mathbf{x}}_0(\mathbf{x}_n, n, \mathcal{G}) = \frac{1}{\sqrt{\bar{\alpha}_n}} (\mathbf{x}_n - \sqrt{1 - \bar{\alpha}_n} \hat{\boldsymbol{\epsilon}}_{\theta}). \quad (27)$$

412 The reverse mean is then computed from the DDPM posterior parame-
 413 terization Ho et al. (2020); Dhariwal and Nichol (2021),

$$\boldsymbol{\mu}_{\theta}(\mathbf{x}_n, n, \mathcal{G}) = \tilde{\boldsymbol{\mu}}(\mathbf{x}_n, \hat{\mathbf{x}}_0) = \frac{1}{\sqrt{\bar{\alpha}_n}} \left(\mathbf{x}_n - \frac{\beta_n}{\sqrt{1 - \bar{\alpha}_n}} \hat{\boldsymbol{\epsilon}}_{\theta} \right). \quad (28)$$

414 For numerical stability, the reconstructed $\hat{\mathbf{x}}_0$ was clipped to the normal-
 415 ized coordinate range $[-1, 1]$ to constrain reconstructed normalized positions
 416 to the training range.

417 *Training objective.* Training sampled a clean trajectory $\mathbf{x}_0$, a diffusion step
 418 n , and Gaussian noise $\boldsymbol{\epsilon}$, then constructed $\mathbf{x}_n$ using Eq. (21). The denoiser
 419 was trained with the standard noise-prediction objective loss function,

$$\mathcal{L}_{\text{ddpm}}(\theta) = \mathbb{E}_{\mathbf{x}_0, n, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_n, n, \mathcal{G})\|_2^2], \quad n \sim \text{Uniform}\{1, \dots, N\}, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (29)$$

Minimizing Eq. (29) trains the U-net ϵ_θ to estimate the sampled noise at diffusion step n , conditioned on the encoded airway geometry. After training, new trajectories were generated by initializing $\mathbf{x}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and iterating over the reverse process from $n = N$ to 1 applying the SDF guidance described in Section 2.5.

Regarding the geometry conditioning, the diffusion model must generate trajectories in bounded anatomy-dependent domains. The occupancy grid therefore supplies a learned global representation of airway morphology to the temporal denoiser. This learned geometry conditioning is distinct from the SDF guidance term, the occupancy grid is encoded by the neural network, whereas the SDF is queried directly during inference as a soft penalty.

2.4.1. Global geometry encoder and variational bottleneck

Given a voxelized occupancy grid $\mathbf{B} \in \{0, 1\}^{1 \times D \times H \times W}$, a three-dimensional variational geometry encoder maps the occupancy volume to a global stochastic embedding. In compact notation,

$$\mathbf{z}_G = f_\phi(\mathbf{B}), \quad \mathbf{z}_G \in \mathbb{R}^{d_e}. \quad (30)$$

The purpose of f_ϕ is to compress the high-dimensional occupancy volume into a low-dimensional vector supplied to the denoiser. The study does not directly evaluate how much local or distal-branch information is retained in this global representation.

The encoder outputs distribution parameters,

$$(\boldsymbol{\mu}_G, \log \boldsymbol{\sigma}_G^2) = f_\phi(\mathbf{B}), \quad \boldsymbol{\mu}_G, \boldsymbol{\sigma}_G \in \mathbb{R}^{d_e}, \quad (31)$$

and samples the embedding through reparameterization,

$$\mathbf{z}_G = \boldsymbol{\mu}_G + \boldsymbol{\sigma}_G \odot \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (32)$$

where $\odot$ denotes elementwise multiplication. The Kullback–Leibler (KL) regularization term was

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(\mathcal{N}(\boldsymbol{\mu}_G, \boldsymbol{\sigma}_G^2) \parallel \mathcal{N}(\mathbf{0}, \mathbf{I})). \quad (33)$$

443 The complete training objective for the conditioned DDPM was

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{ddpm}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}}, \quad \lambda_{\text{KL}} = 10^{-4}. \quad (34)$$

444 The KL term pulls the geometry-embedding distribution toward the prior
 445 and was intended to discourage the encoder from memorizing the discrete
 446 training geometries as isolated codes. It does not, on its own, establish
 447 that anatomically similar airways map to nearby latent points or that the
 448 latent space is organized in an interpretable way. Because the module has
 449 no decoder and no geometry-reconstruction objective, we refer to it as a
 450 variational geometry encoder rather than as a decoder-based reconstruction
 451 model.

452 The trajectory denoiser $\epsilon_{\theta}(\mathbf{x}_n, n, \mathcal{G})$ receives the embedded geometry through
 453 the same embedding pathway used for diffusion-step conditioning. The re-
 454 ported operation adds the geometry embedding to the timestep embedding,

$$\mathbf{e}_{\text{cond}}(n, \mathcal{G}) = \mathbf{e}(n) + \mathbf{z}_{\mathcal{G}}. \quad (35)$$

455 If the timestep and geometry embeddings are not already dimension
 456 matched, a learned projection was done before the sum in (35). The re-
 457 sulting $\mathbf{e}_{\text{cond}}$ conditions each residual block in the temporal U-Net on both
 458 the diffusion step and the encoded airway representation.

459 Following the guided-diffusion U-Net style Nichol and Dhariwal (2021);
 460 Dhariwal and Nichol (2021), each residual block receives an embedding vec-
 461 tor. With $\mathbf{h}_{\ell}$ denoting the feature map at block ℓ , the conditioned residual-
 462 block schema is

$$\mathbf{h}_{\ell+1} = \mathbf{h}_{\ell} + \text{Res}(\mathbf{h}_{\ell}, \mathbf{e}_{\text{cond}}(n, \mathcal{G})). \quad (36)$$

463 and the evaluated implementation injects the embedding as a projected
 464 additive bias inside the residual transformation,

$$\text{Res}(\mathbf{h}_{\ell}, \mathbf{e}_{\text{cond}}) = \text{Conv}_2\left(\sigma\left(\text{GN}\left(\text{Conv}_1\left(\sigma\left(\text{GN}(\mathbf{h}_{\ell})\right)\right) + \mathbf{W}_{\ell}\mathbf{e}_{\text{cond}}\right)\right)\right), \quad (37)$$

465 where $\sigma(\cdot)$ is the SiLU activation, GN is group normalization, and $\mathbf{W}_{\ell}$ is
 466 a learned projection. This is an additive-bias mechanism, not FiLM scale-
 467 and-shift modulation Perez et al. (2018); which we describe in Eq. (37).

468 2.5. SDF-guided diffusion sampling

469 Occupancy conditioning does not guarantee that every stochastic sample
 470 remains within the represented lumen. During inference, an SDF-derived for-
 471 mulation, which can be interpreted as a repulsive energy, was therefore used
 472 to bias reverse diffusion away from out-of-lumen positions. This guidance is
 473 a soft geometric penalty, not a physical wall-contact model and not a hard
 474 containment guarantee, specially since the DDPM is a denoising process,
 475 which does not generate particle tracks one step at a time. The SDF grid
 476 $\mathbf{S} \in \mathbb{R}^{D \times H \times W}$ shares the geometry grid with the occupancy representation,
 477 and $S(\mathbf{x})$ denotes a continuous SDF value queried at a trajectory position.

478 To formulate the energy-like SDF penalty, let $\mathbf{x}_n \in \mathbb{R}^{K \times 3}$ denote the
 479 current noisy trajectory at diffusion step n , with $\mathbf{x}_{n,k} \in \mathbb{R}^3$ the position at
 480 physical-time index k . Using the negative-inside and positive-outside con-
 481 vention from Section 2.3.3, the out-of-lumen energy becomes

$$E(\mathbf{x}_n; \mathbf{S}) = \sum_{k=0}^{K-1} \psi(S(\mathbf{x}_{n,k})). \quad (38)$$

482 The hinge penalty was

$$\psi(s) = \text{ReLU}(s). \quad (39)$$

483 Thus, generated trajectories that have points with $S(\mathbf{x}) < 0$ do not con-
 484 tribute to the hinge penalty, while points with $S(\mathbf{x}) > 0$ contribute in propor-
 485 tion to the interpolated positive SDF value. Away from the non-differentiable
 486 boundary of the hinge penalty, the contribution of a single point is

$$\nabla_{\mathbf{x}_{n,k}} E(\mathbf{x}_n; \mathbf{S}) = \begin{cases} \nabla_{\mathbf{x}} S(\mathbf{x}_{n,k}), & S(\mathbf{x}_{n,k}) > 0, \\ \mathbf{0}, & S(\mathbf{x}_{n,k}) < 0. \end{cases} \quad (40)$$

487 The descent direction $-\nabla_{\mathbf{x}_{n,k}} E$ is intended to move positive-SDF points
 488 toward lower SDF values and the represented lumen. Because the SDF is
 489 interpolated on an approximate voxel grid, this operation is as an inference-
 490 time bias rather than a guarantee that all samples remain in-lumen.

491 2.5.1. Guidance-modified reverse step

492 In standard DDPM sampling, one reverse step is

$$p_{\theta}(\mathbf{x}_{n-1} \mid \mathbf{x}_n, \mathcal{G}) = \mathcal{N}(\boldsymbol{\mu}_{\theta}(\mathbf{x}_n, n; \mathcal{G}), \boldsymbol{\Sigma}_n), \quad (41)$$

where μ_θ is computed from the denoiser as described in equation 28. The reported guided sampler shifted the reverse mean using the SDF energy descent direction,

$$\tilde{\mu}(\mathbf{x}_n, n; \mathcal{G}, \mathbf{S}) = \mu_\theta(\mathbf{x}_n, n; \mathcal{G}) - \lambda_g \nabla_{\mathbf{x}_n} E(\mathbf{x}_n; \mathbf{S}), \quad (42)$$

with guidance scale $\lambda_g > 0$. The value of $\lambda_g = 1.0$ was set as the default value, and in the experiments ranged from 1.0 to 4.0, adjusted on the fly to reduce the amount of trajectories that drifted away from the airway lumen. The guided reverse sample was then

$$\mathbf{x}_{n-1} = \tilde{\mu}(\mathbf{x}_n, n; \mathcal{G}, \mathbf{S}) + \Sigma_n^{1/2} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (43)$$

This mean shift is described as the implemented SDF-guided sampling rule. It should not be read as a claim that the method implements covariance-scaled classifier Dhariwal and Nichol (2021).

2.5.2. Sampling algorithm

Algorithm 1 summarizes the geometry-conditioned sampling procedure. The occupancy grid $\mathbf{B}$ conditions the denoiser throughout the reverse process, while the SDF $\mathbf{S}$ contributes through the gradient term in Eq. (42). Physical-world coordinates of the trajectories are obtained by de-normalizing $\mathbf{x}_0$ after sampling.

2.6. Network architecture and training protocol

The last part related to the DDPM is the actual model architecture. The denoiser was a one-dimensional U-Net operating along the trajectory-time dimension. Its input shape was (B, C, K) , with $C = 3$ position channels and $K = 1024$ time indices, and its output was $\hat{\epsilon} \in \mathbb{R}^{B \times C \times K}$. The U-Net followed the guided-diffusion design Nichol and Dhariwal (2021); Dhariwal and Nichol (2021), with base channel count 128, channel multipliers (1, 1, 2, 3, 4), and 4 attention heads.

Training used AdamW with an initial learning rate of 10^{-4} , batch size 128 or 64 depending on the GPU configuration, and an exponential moving average (EMA) of model weights. Multi-GPU training used MPI/DDP following the implementation pattern in Li et al. (2024a). The available hardware was either one NVIDIA H100 GPU or two NVIDIA A100 GPUs, each with 80 GB of memory. The GitHub repository provides all the implementation and configuration details.

Algorithm 1 Geometry-conditioned diffusion sampling with SDF-based soft guidance

Require: Occupancy grid $\mathbf{B}$ representing $\mathcal{G}$, SDF $\mathbf{S}$, metadata $(\mathbf{o}, \mathbf{h}, \mathbf{d})$, guidance scale λ_g , diffusion steps N

- 1: Initialize $\mathbf{x}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in normalized coordinate space
- 2: **for** $n = N, N - 1, \dots, 1$ **do**
- 3: Predict noise: $\hat{\epsilon} \leftarrow \epsilon_\theta(\mathbf{x}_n, n; \mathbf{B})$
- 4: Compute DDPM mean: $\mu_\theta(\mathbf{x}_n, n; \mathbf{B})$ from $\hat{\epsilon}$
- 5: Evaluate SDF energy: $E \leftarrow \sum_k \text{ReLU}(S(\mathbf{x}_{n,k}))$
- 6: Compute soft-guidance direction: $\mathbf{g} \leftarrow -\nabla_{\mathbf{x}_n} E$
- 7: Shift mean: $\tilde{\mu} \leftarrow \mu_\theta + \lambda_g \mathbf{g}$
- 8: Sample: $\mathbf{x}_{n-1} \leftarrow \tilde{\mu} + \Sigma_n^{1/2} \mathbf{z}$, $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 9: **end for**
- 10: De-normalize $\mathbf{x}_0$ to world coordinates when reporting trajectories in physical units
- 11: **return** $\mathbf{x}_0$

Figure 2 provides a graphical summary of the training and guided-inference workflows. The upper row presents the three principal components of training. First, the voxelized airway geometry is mapped to a stochastic anatomy embedding using a variational encoder. In parallel, a clean particle trajectory is progressively corrupted through the DDPM forward process according to a prescribed noise schedule, such that the trajectory distribution approaches a standard Gaussian distribution at the terminal diffusion step. The resulting noisy trajectory, diffusion-step embedding, and geometry latent are then provided to the one-dimensional U-Net, which is trained to predict the injected noise at randomly sampled diffusion steps. Model parameters are optimized by minimizing the discrepancy between the predicted and applied noise.

The lower row illustrates the guided inference procedure. Sampling begins from Gaussian noise, and the airway geometry is again encoded to condition the reverse denoising process. Conditional and unconditional noise estimates are combined through classifier-free guidance, while an additional SDF-based guidance term penalizes trajectory points predicted outside the airway lumen. This soft geometric constraint biases the reverse process toward anatomically feasible particle tracks, although it does not impose an exact physical boundary condition or explicitly model particle-wall interactions.




Figure 2: **Methodological overview of the geometry-conditioned guided DDPM.** The top row summarizes training. A subject-specific airway geometry is voxelized and encoded with a variational geometry encoder to obtain a stochastic global anatomy embedding. A clean fixed-length particle-position trajectory is progressively corrupted by the forward diffusion process, and a one-dimensional U-Net denoiser is trained to predict the injected noise while conditioned on the diffusion step and encoded geometry. The bottom row summarizes inference. Starting from Gaussian noise, the reverse process generates a position trajectory. The occupancy grid provides learned anatomical conditioning, while the SDF supplies an inference-time soft guidance term intended to reduce out-of-lumen samples. The guidance term does not enforce an exact physical boundary or predict terminal deposition or escape states.

2.7. Evaluation design and trajectory-derived metrics

Generated trajectories were compared with the CFPD simulation-derived reference trajectories using the existing qualitative and quantitative summaries which will be described in detail in this section. The qualitative comparisons use projected and three-dimensional overlays of reference and generated tracks over the airway geometry. The quantitative comparisons use speed magnitude distributions, path-length and acceleration summaries, Wasserstein-1 distance for pooled speed distributions, and, a further analysis was used comparing DDPM to DMD reduced order model outputs. Table 1 reports the trajectory-derived kinematic summaries and speed-distribution distances for the main DDPM evaluation.

The definitions below use the position trajectory $\{\mathbf{x}_k\}_{k=0}^{K-1}$ from Eq. (16). For both reference and generated tracks, finite-difference quantities were computed from positions on the fixed grid with time step Δt . Metrics pooled

558 over all time indices excluded stationary terminal segments introduced by
 559 last-value hold. If included, these segments add zero or near-zero increments
 560 after a track terminates, affecting the interpretation of pooled speed and
 561 acceleration distributions.

562 2.7.1. Kinematic metrics

563 The velocity was computed as forward finite-difference associated with
 564 the interval from k to $k + 1$ as,

$$\hat{\mathbf{v}}_k = \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\Delta t}, \quad k = 0, \dots, K - 2.$$

565 The speed magnitude was defined as

$$s_k = \|\hat{\mathbf{v}}_k\|_2,$$

566 where $\|\cdot\|_2$ is the Euclidean norm. Speed has units of m s^{-1} . The reported
 567 speed summaries include the mean $\mathbb{E}[s]$, standard deviation $\text{Std}[s]$, and 99th
 568 percentile $\text{p99}[s]$.

569 Acceleration magnitude was approximated by a first-order finite difference
 570 of the velocities,

$$a_k = \frac{\|\hat{\mathbf{v}}_{k+1} - \hat{\mathbf{v}}_k\|_2}{\Delta t}, \quad k = 0, \dots, K - 3.$$

571 Acceleration is reported in units of m s^{-2} . The reported acceleration sum-
 572 maries include the mean $\mathbb{E}[a]$ and 99th percentile $\text{p99}[a]$.

573 Another metric analyzed was the path length for a trajectory; which was
 574 computed from position increments,

$$L = \sum_{k=0}^{K-2} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_2. \quad (44)$$

575 Path length has units of meters.

576 2.7.2. Distributional similarity

577 The one-dimensional Wasserstein-1 distance was used to compare em-
 578 pirical pooled speed distributions for generated and reference trajectories
 579 Gabriel and Marco (2019). We define F_{gen} and F_{ref} to denote the empiri-
 580 cal cumulative distribution functions of the generated and reference speeds.
 581 Then

$$W_1(F_{\text{gen}}, F_{\text{ref}}) = \int_{-\infty}^{\infty} |F_{\text{gen}}(u) - F_{\text{ref}}(u)| du. \quad (45)$$

582 For speed distributions, W_1 has units of m s^{-1} . While there is no strict lower
 583 bound, lower values indicate closer pooled speed distributions under this
 584 metric.

585 *2.8. Exploratory DMD baseline comparison*

586 The DMD comparison from previous work by Graffigna et al. (2026) was
 587 retained as an exploratory baseline focused on particle-track representation.
 588 DMD is a deterministic, geometry-specific reconstruction fitted to a fixed
 589 ensemble state, whereas the DDPM developed in this work is a conditional
 590 stochastic generator of subject-specific fixed-length trajectory samples. DMD
 591 and multiresolution extensions are described in the cited literature SCHMID
 592 (2010); ROWLEY et al. (2009); Kutz et al. (2016); the present manuscript
 593 uses them only as a contrast for selected trajectory diagnostics. The com-
 594 parison was restricted to a single geometry, namely Geometry 1, for which
 595 the DDPM was trained and the DMD operator was fitted.

596 The DMD reconstructed particle trajectories was splitted in two tensors,
 597 one with an all-particle tracks in a single tensor (both stuck and particles that
 598 have scaped) and a “stuck,” tensor, retaining only particles that were stuck to
 599 the lumen wall. Post hoc jump filtering for DMD was applied before comput-
 600 ing derivative-sensitive diagnostics. A step exceeding $5 \times 10^{-2} \text{ m}$ was flagged.
 601 Trajectories with multiple flagged jumps were excluded, and trajectories with
 602 one flagged jump were truncated at the first flagged index. Tracks with fewer
 603 than 55 valid frames were excluded. The filtered comparison retained 404
 604 all-tracks trajectories and 290 stuck trajectories, for 694 total retained DMD
 605 trajectories. The 404 retained all-tracks trajectories corresponded to 2.3%
 606 of the reported 17,840 all-tracks DMD output. This filtering reduced the
 607 influence of discontinuities on finite-difference velocities and spectra, but it
 608 also biased the DMD sample toward the most stable reconstructions.

609 The DMD exports contained 60 and 55 frames for the two groups. DDPM
 610 and reference trajectories were sampled on a common 55-frame grid using
 611 indices selected from the valid 0–1023 range of the 1024-element sequence.
 612 Under the adopted effective mapping, $\Delta t \approx 2.370 \times 10^{-2} \text{ s}$. Because the
 613 DMD and DDPM/reference tensors were produced through different time
 614 pipelines, lag and spectral diagnostics are best interpreted as frame-matched
 615 comparisons. When seconds or hertz are reported for this comparison, they
 616 refer to the adopted effective mapping.

617 The DMD diagnostics included scalar marginal comparisons, component-
 618 velocity distributions, multi-lag increment statistics, autocorrelation, and

619 power spectral density. For scalar marginals, the reported analysis used
620 Wasserstein-1 distance, the Kolmogorov-Smirnov statistic, and energy distance
621 Ramdas et al. (2017); Massey (1951); Székely and Rizzo (2013). The
622 corresponding filtered comparison is reported in Figure 7, Figure 11, Figure
623 12, and Table 2. Any reported comparison between DMD and guided
624 DDPM should therefore be read as an exploratory, filtered, single geometry
625 contrast rather than as a definitive ranking of the two modeling approaches.

626 3. Results

627 We evaluated the geometry-conditioned guided DDPM by comparing its
628 generated three-dimensional particle-position trajectories against simulation-
629 derived reference trajectories. Four questions organize the section: how SDF
630 guidance affects apparent trajectory containment and airway-topology alignment;
631 how closely generated tracks agree with the reference on training and
632 held-out airway geometries; how well selected trajectory-derived kinematic
633 distributions and summary metrics match; and how the guided DDPM compares,
634 with a DMD particle-track reconstruction, which is another data-driven
635 reduced order model. Deposited dose, regional deposition, escape probability,
636 and categorical terminal-event prediction were not evaluated.

637 A note on sample counts for the figures, for visual clarity, the trajectory-
638 overlay figures display 256 randomly selected trajectories out of the 16,384
639 generated for each set, whereas the distributional summaries use the full
640 reported counts rather than the displayed subset. In the multi-geometry
641 velocity analysis, the source reports 16,384 generated and 50,000 reference
642 trajectories, and the DMD comparison later reports 16,384 trajectories of
643 1,024 timesteps, only for a single geometry.

644 3.1. Qualitative effect of SDF guidance on generated trajectories

645 Figure 3 shows reference trajectories, unguided DDPM samples, and SDF-
646 guided DDPM samples for two reported airway geometries. In the unguided
647 samples, some trajectory segments visibly extend outside the displayed air-
648 way outline, and several tracks do not follow the reference bifurcation pattern
649 in the projected views. These visual discrepancies are clear in the tracheal
650 and main bronchial regions, where the unguided samples drift laterally relative
651 to the represented lumen.

652 On the other hand, the SDF-guided trajectories display lower visually
653 apparent out-of-lumen positions and closer alignment with the represented




Figure 3: **Qualitative effect of SDF guidance on generated particle trajectories.** Reference trajectories (green), unguided DDPM samples (orange), and SDF-guided DDPM samples (blue) are shown for Geometries 6 and 7. Each geometry is displayed in the x - y , x - z , and y - z projections and in a three-dimensional view; the airway surface outline is shown in gray. Guided samples show fewer visually apparent out-of-lumen segments in these examples, but SDF acts as a soft bias and does not guarantee complete lumen containment.

major branches. While this is an improvement, SDF guidance was used as a soft inference-time geometric bias; it does not model wall-interaction physics, therefore there are still some boundary inconsistencies.

3.2. Trajectory generation on training and held-out airway geometries

Figure 4 summarizes qualitative trajectory agreement across five reported geometries. Geometries 1 and 2 are described as training geometries, whereas Geometries 3, 6, and 7 are described as held-out geometries. For the training geometries, the generated trajectories follow the displayed major airway pathways and branch-wise splitting patterns in the projections. Differences remain visible in smaller distal branches, where generated trajectories appear smoother or less spatially dispersed than the reference trajectories. For the reported held-out geometries, the geometry-conditioned sampling framework generated trajectories that are in agreement with the global airway topology, including the main airway and primary bifurcations. Discrepancies increased in distal branches, where generated tracks may under-represent peripheral pathways, exhibit lateral offsets relative to the surface outline, or concentrate in one dominant branch compared with the reference ensemble. These findings provide preliminary evidence of anatomy-conditioned

trajectory generation on the stated held-out cases and establish potential for broader anatomical robustness.

Two sources of test-geometry information are present in the reported framework, first the voxelized occupancy grid is passed through the learned variational geometry encoder, which provides a global conditioning embedding. Separately, the SDF for the same geometry is supplied directly during reverse diffusion as an inference-time soft geometric guide. Therefore, held-out topology alignment and apparent containment are not solely attributed to the learned geometry embedding.

Due to the cost of CFPD simulations, only nine simulated airway geometries had reference particle trajectories used for training/testing the generative model. The qualitative held-out results should consequently be interpreted within the limited anatomy set. Additional simulation-derived trajectories from a more diverse anatomy set would be useful for future evaluation.

3.3. Velocity magnitude and trajectory-derived statistics

Velocity magnitudes were derived from first-order finite differences of generated and reference position trajectories. For the multi-geometry velocity comparison, the reported sample counts were $N = 16,384$ generated trajectories and $N = 50,000$ reference trajectories. Figures 5 and 6 show cumulative distribution functions and probability-density histograms for velocity magnitude across the analyzed geometries. Geometry 8 was excluded from this analysis because there was a large mismatch between the generated particle tracks by the DDPM and the represented geometry, despite agreement in part of the central velocity range.

Across the displayed geometries, the generated and reference cumulative distributions agree most closely in the central portion of the speed range. The histograms also show similar distributional shapes in several cases, while residual differences remain near the zero-velocity component and in the high-speed tails. The use of diffusion models for trajectory-sampling problems is consistent with the broader modeling context discussed by Li et al. (2024a), and the present evidence is limited to the reported airway trajectory validation.

Regarding quantitative analysis, percentage differences for the evaluation set with $K = 1024$ (timesteps) and $\Delta t = 0.00125$ were computed. The DDPM generated mean speed differed by approximately -4.3% , and the 99th percentile speed differed by approximately -12.0% . Similar reported differ-




Figure 4: **Reference and guided DDPM trajectories across reported airway geometries.** Five airway geometries are shown with particle trajectories projected onto the x - y , x - z , and y - z planes and in a three-dimensional view. Reference trajectories are shown in green, guided DDPM samples are shown in blue, and the airway surface outline is shown in gray. Geometries 1 and 2 are reported as training geometries, whereas Geometries 3, 6, and 7 are reported as held out. The figure provides a qualitative comparison on the stated split and should not be read as proof of broad generalization.

ences were given for mean path length (-4.3%), mean acceleration (-3.8%),
and acceleration 99th percentile (-11.5%).

Taken these results together, the qualitative trajectory overlays and velocity-




Figure 5: **Velocity-magnitude distributions for Geometries 1–4.** Cumulative distribution functions (left) and probability-density histograms (right) compare generated trajectories with reference trajectories. Velocity magnitudes were derived from position differences. Last-value-held terminal segments contribute near-zero mass, so the low-speed region should be interpreted together with the terminal-padding representation described in the Methods.




Figure 6: **Velocity-magnitude distributions for Geometries 5–7.** Cumulative distribution functions (left) and probability-density histograms (right) compare generated trajectories with reference trajectories. Residual discrepancies are visible near the stationary component and in the upper tail for some geometries. Velocity magnitudes were computed from position differences, and last-value-held terminal segments contribute to the near-zero component.

711 distribution comparisons show that the guided DDPM generated position
 712 trajectories that followed the reported airway topology and reproduced tra-
 713 jectory kinematics to varying degrees. The evidence remains limited by the
 714 small number of simulated geometries, by the position-only trajectory rep-
 715 resentation, and by the absence of direct evaluation of terminal outcomes.
 716 The results should therefore be interpreted in the light of a proof-of-concept

Table 1: **Summary of trajectory-derived kinematic metrics and speed-distribution distance.** For each listed geometry, reference and guided DDPM values are reported for mean acceleration magnitude, 99th percentile acceleration magnitude, mean path length, mean speed magnitude, 99th percentile speed magnitude, speed standard deviation, and Wasserstein-1 distance between pooled speed distributions. Acceleration is in m s^{-2} , path length is in m, and speed quantities and speed W_1 are in m s^{-1} . These metrics characterize generated position sequences and their finite-difference kinematics; they do not assess deposition outcomes.

Geometry	Set	Accel. mean (m s^{-2})	Accel. p99 (m s^{-2})	Path length (m)	Speed mean (m s^{-1})	Speed p99 (m s^{-1})	Speed std (m s^{-1})	Speed W_1 (m s^{-1})
1	Reference	55.33	568.3	0.2822	0.2207	1.680	0.4291	–
	Guided DDPM	54.77	518.9	0.2798	0.2188	1.528	0.4042	0.0181
2	Reference	54.86	518.1	0.2781	0.2175	1.520	0.3967	–
	Guided DDPM	54.92	505.1	0.2761	0.2159	1.478	0.3877	0.00618
3	Reference	64.73	561.1	0.3113	0.2435	1.653	0.4578	–
	Guided DDPM	60.33	524.3	0.3079	0.2408	1.550	0.4034	0.0209
4	Reference	45.44	535.3	0.2273	0.1778	1.599	0.3853	–
	Guided DDPM	45.01	531.8	0.2259	0.1767	1.597	0.3837	0.00135
6	Reference	58.44	550.7	0.2995	0.2342	1.602	0.4247	–
	Guided DDPM	53.95	532.0	0.2717	0.2125	1.546	0.4028	0.0217
7	Reference	55.15	543.4	0.2826	0.2210	1.604	0.4171	–
	Guided DDPM	53.68	523.0	0.2773	0.2168	1.528	0.4053	0.00509

trajectory-generation evaluation.

3.4. Quantitative trajectory metrics

Table 1 reports acceleration, path-length, speed, and speed-distribution values for the analyzed HRT geometries. The metric definitions are provided in Section 2.7. The reported W_1 column compares pooled speed distributions between generated and reference trajectories, with units of m s^{-1} . Lumen containment was over 95% for all geometries, and the reported metrics were computed on the retained trajectories including out-of-lumen segments.

The table includes Geometries 1, 2, 3, 4, 6, and 7. Based on the split stated in the trajectory-overlay results, Geometries 1, 2, and 3 were seen by the encoder during training and Geometries 4, 6, and 7 are held-out cases.

Across the reported geometries, mean values were generally closer than extreme percentiles, while the magnitude of the discrepancy varied by anatomy. The speed W_1 values also varied across cases, with lower values indicating closer pooled speed distributions relative to the reference for each of the geometries.

733 3.5. Exploratory comparison with DMD for Geometry 1

734 3.5.1. Evaluation setup

735 This comparison is an exploratory contrast between two data-driven ap-
736 proaches. The guided DDPM is a conditional stochastic generator of in-
737 dividual particle-position trajectories, using a learned geometry embedding
738 and SDF guidance during sampling. The DMD reconstruction on the other
739 hand is a deterministic, single-geometry reduced-order reconstruction fitted
740 to a fixed particle-ensemble state. The results below should not be inter-
741 preted as a general benchmark of surrogate quality or as evidence that one
742 approach replaces the other, but a comparison of their respective strengths
743 and limitations.

744 The DMD particle-track reconstruction was taken from the DMD study
745 Graffigna et al. (2026). In that formulation, the three Cartesian coordinates
746 of all particles at a given time were shifted by injection time, $\tilde{t} = t - t_{\text{injected}}$,
747 and stacked into a single state tensor. A reduced linear operator was then
748 fitted to advance the ensemble state in time. The Koopman operator identi-
749 fies an approximate linear evolution operator from paired snapshot matrices
750 $\mathbf{X}^- = [\mathbf{x}_1, \dots, \mathbf{x}_{m-1}]$ and $\mathbf{X}^+ = [\mathbf{x}_2, \dots, \mathbf{x}_m]$, so that $\mathbf{x}_{k+1} \approx \mathbf{A}\mathbf{x}_k$; the eigen-
751 values and modes of $\mathbf{A}$ define exponential space-time components SCHMID
752 (2010); ROWLEY et al. (2009). Multiresolution DMD extends this idea by
753 recursively applying DMD over time windows to separate dynamics across
754 temporal scales Kutz et al. (2016). Here, DMD was used only as a particle-
755 track reconstruction baseline for the reported Geometry 1 comparison.

756 The comparison was restricted to Geometry 1 because the DMD operator
757 was fitted in one anatomical domain and cannot be applied unchanged to a
758 new airway mesh without generating new snapshots, mapping states between
759 geometries, and re-identifying the Koopman operator. The guided DDPM
760 samples and the DMD reconstruction were therefore compared on the same
761 geometry. The comparison focuses on particle-track representation, spatial
762 continuity, and same kinematic distributions. The DMD reconstruction was
763 fitted to two exported tensors: one constructed from particles stuck at the
764 wall and another including all particles that either escaped the domain or
765 terminated at the wall. These event categories describe the source trajectory
766 exports and DMD inputs; the guided DDPM results reported here remain
767 position-sequence outputs and do not include an explicit terminal-event clas-
768 sifier.

769 Figure 7 gives the qualitative Geometry 1 comparison. The guided DDPM




Figure 7: **Exploratory Geometry 1 track comparison between guided DDPM and DMD.** The black curves denote the STL outline, green curves denote the DMD all-tracks group, red curves denote the DMD stuck-particle group, and blue curves denote guided DDPM samples. Both outputs occupy major portions of the primary airway envelope, but the unfiltered DMD output includes isolated discontinuous return-to-origin segments. The derivative-based DMD diagnostics below use the jump-filtered subset described in the text and should be interpreted as a filtered Geometry 1 comparison.

770 samples occupy the primary airway geometry in the X - Y , X - Z , Y - Z , and
 771 three-dimensional views. The DMD reconstruction also recovers major por-
 772 tions of the displayed branches, but some DMD trajectories contain abrupt
 773 non-physical discontinuous segments that return toward the coordinate origin
 774 near the carina. These discontinuities are reconstruction artifacts that come
 775 from the DMD fitting procedure, and they are not present in the reference
 776 trajectories. The DMD tracks artifacts were filtered out before derivative-
 777 based metrics were computed, as described below.

778 The DMD tracks were filtered before derivative-based metrics were com-
 779 puted. A particle track step was flagged when its magnitude exceeded
 780 5×10^{-2} m, which is as approximately three orders of magnitude above a
 781 typical high-velocity region local step in the clean trajectories. Particles
 782 with multiple flagged jumps were excluded because repeated discontinuities
 783 would dominate velocities and power spectra. For trajectories with a single
 784 flagged jump, the first flagged index was interpreted as the onset of non-
 785 physical return-to-origin decay, and subsequent samples were discarded. Af-
 786 ter filtering, particles with at least $K = 55$ valid pre-decay timesteps were
 787 retained. The resulting filtered DMD set contained 404 all-tracks particles
 788 and 290 stuck-particle trajectories, for 694 retained trajectories. This fil-
 789 tering reduced the influence of discontinuous artifacts on derivative-based
 790 summaries, but it also biased the retained DMD set toward the most stable
 791 reconstructions.

792 The DMD reconstructions spanned 60 timesteps for the all-tracks group

and 55 timesteps for the stuck-particle group. The diffusion samples were uniformly subsampled from $K_{\text{full}} = 1024$ to $K = 55$, with the same frame selection applied to the reference tracks. Under this mapping, the common frame spacing was taken as $\Delta t \approx 1.279 \text{ s}/54 \approx 2.370 \times 10^{-2} \text{ s}$. Because the DMD export and the DDPM/reference tensor were generated through different pipelines, temporal diagnostics that depend on an absolute time axis, including autocorrelation lag and PSD frequency, should be interpreted primarily as frame-matched comparisons.

For speed, velocity components, and multi-lag increment norm, the reported comparison used Wasserstein-1 distance W_1 , the Kolmogorov–Smirnov statistic D_{KS} , and energy distance Ramdas et al. (2017); Massey (1951); Székely and Rizzo (2013). The definitions and interpretation of these quantities were summarized in Methods section. Table 2 emphasizes W_1 and retains D_{KS} as a secondary diagnostic. The KS p-values were reported as <0.05 .

3.5.2. Speed and component velocities

Speed and component velocities were computed from first-order position differences, $\hat{\mathbf{v}}_k = (\mathbf{x}_{k+1} - \mathbf{x}_k)/\Delta t$, on the common filtered Geometry 1 frame grid. The mean speed across all particles and timesteps was 0.216 m s^{-1} for the reference trajectories, 0.222 m s^{-1} for guided DDPM, and 0.214 m s^{-1} for DMD. The DMD mean was close to the reference mean, but the distribution differed. Speed samples below 0.01 m s^{-1} accounted for 67.8% of DMD samples, compared with 72.5% for guided DDPM and 73.3% for the reference trajectories. The DMD 99th percentile, 1.94 m s^{-1} , exceeded the reference value (1.32 m s^{-1}) and the guided DDPM value (1.36 m s^{-1}) by approximately 45%. The reported speed W_1 values were $6.19 \times 10^{-2} \text{ m s}^{-1}$ for DMD and $6.59 \times 10^{-3} \text{ m s}^{-1}$ for guided DDPM, corresponding to a $9.4\times$ lower speed-distribution distance for guided DDPM.

Figures 8 and 9 show the corresponding distributions (note that the oversaturated near-zero bin was omitted here and in the analysis). The guided DDPM and reference distributions have similar bulk locations, with a peak near $|v| \approx 0.7 \text{ m s}^{-1}$ and tapering above $|v| \approx 1.4 \text{ m s}^{-1}$. The DMD density distribution contains higher probability in the low-speed interval $|v| \in [0.0, 0.4] \text{ m s}^{-1}$ and a slower-decaying right tail. The empirical cumulative distributions show the same redistribution: the DMD curve lies to the left of the reference through part of the bulk distribution and then crosses into a heavier high-speed tail. Thus, confirming that similar mean speeds do not imply similar speed distributions.




Figure 8: Speed distribution $p(|v|)$ for guided DDPM (blue), filtered DMD (orange), and reference trajectories (green) in Geometry 1. The first histogram bin, which contains the near-zero stationary-particle plateau caused in part by last-value-held terminal segments, is suppressed so that the distribution above the plateau is visible. DMD has more mass in the plotted low-speed interval and a heavier upper tail in this filtered comparison.

830 3.5.3. Increment statistics across lags

831 A lag τ denotes the number of frames separating two particle positions.
832 The multi-lag increment norm, $\|\Delta \mathbf{x}(\tau)\|_2 = \|\mathbf{x}_{k+\tau} - \mathbf{x}_k\|_2$, measures the
833 distance traveled over that frame lag. Short lags compare local step sizes,
834 whereas longer lags compare accumulated branch-scale displacement and the
835 temporal ordering of motion along the trajectory. This metric is useful be-
836 cause two trajectory sets can have similar instantaneous speed distributions
837 while differing in how those speeds are arranged along each track.

838 The reported analysis used $\tau \in \{1, 4, 16, 32\}$. Under the adopted effective
839 time mapping, these lags correspond to approximately 0.024 s to 0.76 s, but
840 the primary interpretation remains frame matched. Figure 10 reports the
841 increment-norm densities (note, again the first bin suppressed). Across the
842 four lags, the guided DDPM distributions are closer to the reference distri-
843 butions than the filtered DMD distributions by the reported W_1 values in
844 Table 2. The filtered DMD distributions develop narrow peaks at character-




Figure 9: Empirical cumulative distribution functions of $|v|$ for the filtered Geometry 1 comparison, with samples below 0.01 m s^{-1} omitted to show the dynamic range above the stationary component. The curves characterize pooled finite-difference speeds and should not be interpreted as terminal-event or deposition metrics.

istic magnitudes, including approximately 0.04 m and 0.16 m , that are not apparent in the reference distributions. These peaks become more prominent at longer lags, consistent with repeated modal excursions in the filtered DMD reconstruction.

3.5.4. Temporal correlations and spectra

The normalized speed autocorrelation, $R_{|v|}(\tau)$, measures how speed at one frame is correlated with speed at a later frame. It was included because the speed histogram alone does not determine whether speeds are placed in a similar temporal order along each particle path. For each trajectory, the reported analysis mean-centered the speed sequence, correlated it with a lagged copy, zero-offset normalized the result so that $R_{|v|}(0) = 1$, and averaged over trajectories.

Figure 11 shows similar guided DDPM and reference autocorrelation curves in the source analysis. Under the adopted effective time mapping, both curves decay from 1 toward zero by approximately $\tau \approx 0.4 \text{ s}$, reach a mini-




Figure 10: Increment-norm densities $p(\|\Delta \mathbf{x}(\tau)\|_2)$ at four frame lags for guided DDPM (blue), filtered DMD (orange), and reference trajectories (green) in Geometry 1. The first bin is suppressed to display the range above the stationary component. The comparison is exploratory because the models represent different tasks and the DMD sample was filtered before derivative-based metrics were computed.

860 mum of about -0.2 near $\tau \approx 0.9$ s, and then recover toward zero. The filtered
 861 DMD curve decays faster, with $R(0.1) \approx 0.5$ compared with approximately
 862 0.7 for the reference, and reaches a shallower trough of approximately -0.13 .
 863 These physical-time values depend on the adopted mapping and should be
 864 interpreted as frame-matched diagnostics rather than independently verified
 865 physical lag estimates.

866 The same temporal structure was evaluated spectrally using the one-sided
 867 power spectral density (PSD) of each velocity component. PSD decomposes
 868 the variance of a time series across frequency, and its relationship to the
 869 autocovariance function follows the Wiener–Khinchin theorem Zbilut and
 870 Marwan (2008). In this frame-matched analysis, lower frequencies corre-
 871 spond to slowly varying component motion along coherent airway paths,




Figure 11: Autocorrelation of speed $R_{|v|}(\tau)$ for guided DDPM (blue), filtered DMD (orange), and reference trajectories (green) in Geometry 1. Lags are primarily interpreted in frames; conversions to seconds use the adopted effective mapping described in the text.

872 whereas higher frequencies correspond to rapid component-to-component
873 fluctuations.

874 Figure 12 plots the mean PSD on log-log axes. The reference and guided
875 DDPM spectra follow a decaying shape, with the reported PSD dropping by
876 roughly two decades between $f \approx 1$ Hz and the adopted Nyquist frequency
877 $f_{Ny} \approx 21$ Hz. The guided DDPM spectra were reported within approximately
878 a factor of 1.5 of the reference across the plotted band for v_x and v_y , and
879 within a factor of 2 for v_z . The filtered DMD spectrum differs by remaining
880 flatter or rising with frequency, most visibly for v_x . At the adopted $f \approx 10$ Hz
881 value, the DMD PSD for v_x exceeded the reference by about an order of
882 magnitude, while near $f \approx 1$ Hz it was roughly an order of magnitude lower.

883 3.5.5. Summary of distributional metrics

884 Table 2 consolidates the reported Wasserstein distances and KS statistics
885 for the metrics considered in this analysis. Across the selected diagnostics,
886 the guided DDPM had lower reported W_1 values than the filtered DMD
887 reconstruction. The DMD/guided ratios ranged from 2.1 for v_z to 16.0 for the




Figure 12: Mean one-sided PSD of each velocity component for guided DDPM (blue), filtered DMD (orange), and reference trajectories (green) in Geometry 1. Frequency values in hertz use the adopted effective frame-to-time mapping, so the spectral comparison should be interpreted primarily as a frame-matched diagnostic.

888 $\tau = 16$ increment norm. Because the models represent different tasks and the
 889 DMD output was heavily filtered, this comparison should be interpreted as
 890 an exploratory contrast rather than a definitive ranking of surrogate models.

891 3.5.6. DMD and DDPM limitations

892 The comparison across Figure 7 to Figure 12 and Table 2 led to a consis-
 893 tent interpretation. Guided DDPM reproduced the ground-truth speed dis-
 894 tribution, component-velocity distributions, multi-lag increment statistics,
 895 autocorrelation, and PSD within the resolution of these diagnostics. DMD
 896 reproduced the gross mean speed scale, but the matching mean masked dis-
 897 tributional differences. DMD combined excess slow-particle probability with
 898 an overly heavy high-speed tail and modal peaks in long-lag displacement.

899 These differences reflected known limitations of DMD reduced-order re-
 900 constructions. First, the learned DMD operator is low-rank and linear in the
 901 selected observable space. This structure is useful for coherent dynamics,
 902 but it is not designed to represent intermittent, particle-specific events such
 903 as wall interactions, branch transitions, sharp geometric turns, or stochas-
 904 tic submicron forcing. Second, the DMD reconstruction is unconstrained by
 905 the airway surface. Unlike the guided DDPM sampler, it does not evaluate
 906 a signed distance function or apply a boundary penalty during generation.
 907 Third, DMD is fitted to a fixed state dimension and coordinate system. A
 908 new airway geometry would require new high-fidelity snapshots and a new
 909 operator, or a validated state-mapping procedure, rather than direct evalu-
 910 ation of the existing DMD model. Fourth, DMD is deterministic after the
 911 operator and amplitudes are specified, so it does not naturally sample the
 912 full conditional distribution of plausible Lagrangian tracks.

Statistic	W_1 guided	W_1 DMD	Ratio DMD/guided	D_{KS} (guided / DMD)
$ v $ [m s ⁻¹]	6.59×10^{-3}	6.19×10^{-2}	9.4	0.732 / 0.732
v_x [m s ⁻¹]	4.40×10^{-3}	2.30×10^{-2}	5.2	0.368 / 0.391
v_y [m s ⁻¹]	4.95×10^{-3}	1.29×10^{-2}	2.6	0.371 / 0.386
v_z [m s ⁻¹]	1.98×10^{-2}	4.21×10^{-2}	2.1	0.375 / 0.387
$\ \Delta \mathbf{x}\ _2$ [m], $\tau = 1$	1.56×10^{-4}	1.47×10^{-3}	9.4	0.732 / 0.732
$\ \Delta \mathbf{x}\ _2$ [m], $\tau = 4$	7.58×10^{-4}	4.12×10^{-3}	5.4	0.717 / 0.717
$\ \Delta \mathbf{x}\ _2$ [m], $\tau = 16$	8.65×10^{-4}	1.38×10^{-2}	16.0	0.629 / 0.629
$\ \Delta \mathbf{x}\ _2$ [m], $\tau = 32$	1.79×10^{-3}	2.41×10^{-2}	13.4	0.381 / 0.381

Table 2: Distributional distances for the exploratory, jump-filtered Geometry 1 comparison. “Guided” refers to the guided DDPM model, and “DMD” refers to the filtered DMD reconstruction. W_1 is the Wasserstein-1 distance, with units determined by the row variable. The ratio is DMD divided by guided DDPM. The KS statistic is reported for completeness, but its values were strongly affected by the near-zero stationary component and should not be interpreted as a standalone practical effect size.

913 The jump-filtering step for the DMD data was deemed fair for numer-
914 ical comparison, but it introduced its own limitation. Removing particles
915 with multiple discontinuities and truncating trajectories at the first return-to-
916 origin event prevented non-physical jumps from dominating finite-difference
917 velocities and PSD values. However, the retained 694-particle DMD ensemble
918 represented only the most stable portion of the reconstruction. The 404
919 retained all-tracks particles corresponded to 2.3% of the 17 840 all-tracks
920 output. The fact that DMD still trailed guided DDPM after this favorable
921 filtering indicated that the remaining discrepancy was not solely an artifact
922 of unfiltered jump events.

923 The PSD comparison in Figure 12 shown that velocity-component spec-
924 tra contained more low-frequency than high-frequency power, reflecting per-
925 sistent translation through the airway with smaller high-frequency fluctua-
926 tions. Guided DDPM reproduced this roll-off. DMD produced a flatter or
927 rising spectrum, indicating underrepresentation of slow coherent motion and
928 overrepresentation of high-frequency oscillatory content. The faster DMD
929 autocorrelation decay in Figure 11 was the corresponding time-domain man-
930 ifestation of the same spectral imbalance.

931 For downstream tasks that depend on single-particle increments (includ-
932 ing deposition-site localization, residence-time analysis, and transport-coefficient
933 estimation) the DMD particle-track surrogate was therefore less reliable than

the guided DDPM model in this comparison. The $\tau = 16$ increment distance, which lies in the lag range relevant to branch-scale transport over the reduced 55-frame window, was $16\times$ larger for DMD than for guided DDPM. The guided DDPM model also retained a practical advantage that DMD in its current version did not: once trained, it could be conditioned on a new airway representation, whereas a fitted DMD operator could not be transferred to a new anatomical mesh without additional simulation and operator identification. On the trained/fitted geometry, the guided DDPM model reproduced single-particle speed, component-velocity, multi-lag increment, autocorrelation, and PSD statistics more accurately than the DMD reconstruction, which one may contradict that this analysis may not be fair due to using a geometry that was on the training dataset, but likewise was for the DMD case.

4. Discussion

Across geometries, the guided denoising diffusion probabilistic model closely matched ground-truth mean kinematics but tends to mildly smooth extreme events. In Table 1, speed mean was typically within a few percent of ground truth, while speed p99 and acceleration p99 were consistently lower than ground truth in most cases, indicating slight under-representation of the highest-speed and highest-acceleration events. This behavior demonstrated consistent with denoising-based generative models where guidance and denoising reduced tail variability when the training distribution was limited.

Wasserstein distances between speed distributions were small (on the order of 10^{-3} to 10^{-2}), indicating that the guided DDPM model recovered the overall speed distribution shape well, even when some tail mass is slightly reduced.

Geometry 4 shown the smallest Wasserstein distance ($W_1 \approx 0.00135$), and the kinematic metrics were extremely close (speed p99 differed negligibly). On the other hand, Geometry 6 stood out due to the mean speed and mean path length being notably lower than ground truth, and W_1 was among the largest, indicating that the generated ensemble departed the most from the ground-truth statistics in this case. This trend was consistent with the qualitative behavior observed in guidance comparisons in Figure 4, for difficult geometries, guidance can improve feasibility but may also bias trajectories toward dominant pathways, effectively having the effect of reducing path length and suppressing high-speed excursions.

970 These results should be interpreted in light of the training regime as well,
 971 although each CFPD simulation yielded many trajectories (rich sampling of
 972 the kinematic distribution), the number of distinct simulated geometries was
 973 relatively small. Diffusion models robustly learn trajectory statistics given
 974 sufficient samples, but geometry-specific generalization is fundamentally con-
 975 strained by the diversity of airway shapes seen during training. Therefore,
 976 the observed performance is encouraging for kinematic representation of the
 977 particle paths and distributional recovery, while the feasibility and gener-
 978 alization limitations point directly to dataset expansion (more geometries)
 979 and improved geometry alignment for signed distance function evaluation as
 980 high-impact next steps.

981 Beyond geometry, additional conditioning variables can be appended in
 982 future ablations, e.g., particle diameter d_p , Stokes number St , inlet flow rate
 983 parameters ($Q_{\max}, T$), or event-type priors (deposition vs escape). These can
 984 enter the model as feature, embeddings added to $\mathbf{e}_{\text{cond}}$ or via feature-wise
 985 linear modulation (FiLM).

986 The qualitative analysis demonstrated that diffusion-based generative
 987 modeling remains a practical and effective paradigm for synthesizing La-
 988 grangian particle trajectories conditioned on subject-specific airway geome-
 989 try. Two aspects of the proposed framework were particularly important and
 990 are detailed below.

991 First, geometry conditioning via a variational autoencoder enabled non-
 992 trivial generalization to geometries not encountered during training. The
 993 number of distinct anatomies available for training and testing was only nine.
 994 This imbalance created a fundamental bottleneck for learning geometry-
 995 specific behaviors, especially in distal branches where inter-subject variability
 996 was higher. In this context, the observed qualitative agreement on unseen
 997 geometries should be interpreted as interpolation within a sparsely sam-
 998 pled geometry manifold rather than broad robustness to arbitrary airway
 999 anatomy. Nevertheless, the ability to produce plausible trajectories on held-
 1000 out cases indicated that the latent geometry representation captured mean-
 1001 ingful anatomical factors for conditioning the diffusion model.

1002 Second, SDF guidance played a critical role during inference by improv-
 1003 ing geometric plausibility. Unguided sampling drifted into non-physical re-
 1004 gions, yielding trajectories that violated airway boundaries or form implausi-
 1005 ble bundles not aligned with the lumen. Incorporating guidance consistently
 1006 suppressed outliers and increased general alignment with airway topology,
 1007 particularly in the trachea and main bronchi. At the same time, the re-

sults suggested a trade-off. Stronger guidance reduced sample diversity by concentrating trajectories in dominant pathways and under-representing peripheral branches, which as a consequence it will demand a larger number of trajectory samples when are synthetically generated, as opposed to CFPD generated ones. This behavior was consistent with a guidance mechanism that biases sampling toward geometric feasibility. In practice, the guidance weight should therefore be treated as a tunable parameter that balances constraint satisfaction against faithful reproduction of trajectory variability.

Another outcome was that the guided DDPM model was able to recover velocity magnitude statistics with good agreement to CFPD-derived results, as evidenced by the cumulative distribution functions and histograms. This finding supports the broader expectation that diffusion models are well-suited to learning complex, multimodal distributions without the training instabilities and mode collapse commonly associated with adversarial objectives Liu et al. (2020); Dhariwal and Nichol (2021). Even under limited geometry diversity, the model benefited from the large number of trajectory samples per geometry and learned robust marginal kinematic statistics.

4.1. Limitations

Several limitations must be emphasized to frame the current contributions appropriately and to guide future work.

The most significant limitation was the small number of geometries with particle-track simulations. With only nine simulated anatomies, the model could not fully learn the breadth of anatomical variability presented in population-scale airway datasets. As a consequence, performance degraded for some of the out-of-distribution geometries, where the model hallucinated trajectory patterns that were not physically meaningful (e.g., incorrect branch volumetric occupation or implausible concentration patterns). This limitation was structural and not easy to solve, since improving generalization will require increasing the number of simulated geometries, which requires extra thousands of CPU-hours.

On the residual boundary violations, although signed distance function guidance improves feasibility, some generated trajectories still crossed physical boundaries in challenging cases (e.g., near complex bifurcations, or in thin distal branches). This indicated that guidance alone on the predicted mean for denoising, was not a complete guarantee of constraint satisfaction, especially if the conditioning latent places the sampler near poorly supported regions of the training distribution. Guidance can reduce non-physical outliers

but may also reduce diversity in distal airways and slightly bias the distribution of rare events. This may be especially relevant for applications where capturing tail behavior (rare high-speed jets, near-wall residence events, or extreme accelerations) is important. Moreover, the current framework conditions primarily on geometry (via the variational autoencoder latent) and used a fixed or implicitly learned particle-size distribution from the training data. Realistic aerosol transport for any material depends strongly on particle diameter and density. Extending the model to incorporate or conditioning particle size (either by providing the diameter as an additional conditioning variable for each trajectory or by modeling mixtures of size-resolved conditional distributions) is a fundamental next step. This would enable inference for arbitrary size distributions and would support patient- and scenario-specific prediction at inference time.

Finally, while diffusion-based inference is fast once trained, expanding the training set with additional geometries or particle distributions requires more high-fidelity CFPD simulations and therefore more computational time. This creates a clear roadmap for improvement, but also highlights the need for careful experimental design and potentially hybrid data strategies (e.g., combining a smaller number of high-fidelity simulations with additional lower-cost simulations, likely one-way coupled, or using synthetic geometry augmentation validated against physics).

4.2. Future work

Future work should extend the present particle-trajectory formulation in several directions. A first task is explicit modeling and evaluation of active, deposited, and escaped terminal states, rather than representing terminal behavior only through position sequences and last-value hold. Such a model would require event-aware training targets and separate validation of terminal-event accuracy.

A second gap to address is conditional control over variables that were fixed or implicit in the current training distribution. Particle diameter, injection state, particle-size distribution, and breathing waveform parameters should be represented as explicit conditioning variables before the model is used for scenario-specific trajectory generation. These additions would need validation against reference trajectories generated under corresponding simulation regimes.

Larger and more diverse anatomy-level datasets would be needed as well for broader validation. Additional simulated geometries would support more

informative train/validation/test splits, more rigorous assessment of anatomy-conditioned generation, and better characterization of failure modes in distal branches. Future model development should also investigate multiscale or local geometry conditioning. A single global occupancy embedding may be insufficient to represent local branch structure, distal airway features, and bifurcation-scale constraints. Local surface, centerline, branch, or patch-based conditioning could be evaluated prospectively, provided that the resulting conditioning variables are defined reproducibly and tested against reference trajectories.

Lastly, additional work is needed to verify SDF and coordinate consistency and to evaluate stronger constrained-sampling methods. Future studies should report quantitative containment diagnostics, guidance-scale sensitivity, and failure cases near thin or high-curvature structures. Direct evaluation of deposition outcomes should be treated as a separate downstream validation step rather than assumed from trajectory-shape or kinematic agreement. Uncertainty, sensitivity, and ablation analyses would further clarify which parts of the framework drive the reported behavior.

5. Conclusions

This work introduced a geometry-conditioned, guided DDPM for generating Lagrangian particle trajectories in subject-specific human airway geometries. Motivated by the need for fast, anatomy-aware particle tracking for aerosol deposition analysis, inhaled drug delivery optimization, and radiation research, the proposed approach addresses critical limitations of CFPD simulations by learning a direct generative surrogate from simulation-derived trajectories. The framework combined a variational autoencoder geometry latent for conditioning with SDF guidance during inference to improve geometric feasibility and reduce non-physical wall-penetrating samples.

Across qualitative trajectory comparisons and quantitative analyses, the guided diffusion model produced physically plausible ensembles that matched key statistical properties of ground-truth trajectories, including on geometries not seen during training. These results are particularly encouraging given the extremely limited number of simulated geometries available. Remaining limitations, including occasional boundary violations and reduced robustness for out-of-distribution anatomies, highlight the importance of expanding the diversity of simulated geometries and enriching conditioning variables (such as particle size and breathing rates) in future studies. Overall,

1118 the proposed method demonstrates a viable pathway toward rapid, geometry-
1119 aware Lagrangian particle trajectory generation, bridging the gap between
1120 high-fidelity simulation and practical subject-specific deployment.

1121 **Funding**

1122 Funding for this project is provided by the Department of Defense -
1123 Peer Reviewed Medical Research Program (DOD PRMRP) under award
1124 number W81XWH-21-1-0984 (SAD, MTR, IB) and National Institutes of
1125 Health/National Institute of Allergy and Infectious Diseases (NIH/NIAID)
1126 Radiation and Nuclear Countermeasures Program under award P01AI165380
1127 (SAD, MTR, MG).

1128 **Acknowledgments**

1129 This research was supported in part through research cyberinfrastructure
1130 resources and services provided by the Partnership for an Advanced Com-
1131 puting Environment (PACE) at the Georgia Institute of Technology, Atlanta,
1132 Georgia, USA.

1133 **Declaration of competing interest**

1134 x The authors of this manuscript declare that they have no known com-
1135 peting financial interests or personal relationships that could have appeared
1136 to influence the work reported in this paper.

1137 **Data Availability**

1138 A permanent repository for CFPD ready geometries is available at `dx.`
1139 `doi.org/10.6084/m9.figshare.24787773` Bartol et al. (2024a). All the other
1140 relevant data are available from the corresponding author upon reasonable
1141 request.

1142 **Code Availability**

1143 The code can be found at `https://github.com/ignaciobartol/airway-lagr-diffusion`

1144 Author Contributions

1145 **Ignacio R. Bartol:** Conceptualization, Formal analysis, Investigation,
1146 Methodology, Software, Validation, Visualization, and Writing - original
1147 draft. **Martin S. Graffigna Palomba:** Investigation, Writing - review
1148 & editing. **Mauricio E. Tano:** Conceptualization, Resources, Supervision,
1149 Validation, Writing - review & editing. **Shaheen A. Dewji:** Conceptual-
1150 ization, Funding acquisition, Project administration, Resources, Supervision,
1151 and Writing - review & editing.

1152 References

- 1153 Bartol, I., Graffigna, M., Retamales, M., Dewji, S., 2024a. 3d stl
1154 models of human airways for cfd and cfpd simulations. URL:
1155 [https://figshare.com/articles/dataset/3D_STL_models_of_](https://figshare.com/articles/dataset/3D_STL_models_of_Human_Airways_for_CFD_and_CFPD_simulations_/24787773)
1156 [Human_Airways_for_CFD_and_CFPD_simulations_/24787773](https://figshare.com/articles/dataset/3D_STL_models_of_Human_Airways_for_CFD_and_CFPD_simulations_/24787773),
1157 doi:10.6084/M9.FIGSHARE.24787773.
- 1158 Bartol, I.R., Graffigna Palomba, M.S., Tano, M.E., Dewji, S.A., 2024b.
1159 Computational multiphysics modeling of radioactive aerosol deposition
1160 in diverse human respiratory tract geometries. Communications Engi-
1161 neering 3. URL: <http://dx.doi.org/10.1038/s44172-024-00296-z>,
1162 doi:10.1038/s44172-024-00296-z.
- 1163 Bartol, I.R., Palomba, M.S.G., Dawson, R.J., Bolch, W.E., Tano, M.E.,
1164 Dewji, S.A., 2025. Subject-specific modeling framework for particle depo-
1165 sition using computational fluid dynamics. Journal of Aerosol Science 190,
1166 106660. URL: <http://dx.doi.org/10.1016/j.jaerosci.2025.106660>,
1167 doi:10.1016/j.jaerosci.2025.106660.
- 1168 Brunton, S.L., Noack, B.R., Koumoutsakos, P., 2020. Machine learning for
1169 fluid mechanics. Annual Review of Fluid Mechanics 52, 477–508. URL:
1170 <http://dx.doi.org/10.1146/annurev-fluid-010719-060214>, doi:10.
1171 1146/annurev-fluid-010719-060214.
- 1172 Darquenne, C., 2014. Aerosol deposition in the human lung in reduced
1173 gravity. Journal of Aerosol Medicine and Pulmonary Drug Delivery 27,
1174 170–177. URL: <http://dx.doi.org/10.1089/jamp.2013.1079>, doi:10.
1175 1089/jamp.2013.1079.

- 1176 Dhariwal, P., Nichol, A., 2021. Diffusion models beat gans on
1177 image synthesis, in: Ranzato, M., Beygelzimer, A., Dauphin,
1178 Y., Liang, P., Vaughan, J.W. (Eds.), Advances in Neural In-
1179 formation Processing Systems, Curran Associates, Inc.. pp. 8780–
1180 8794. URL: [https://proceedings.neurips.cc/paper_files/paper/](https://proceedings.neurips.cc/paper_files/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf)
1181 [2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf).
- 1182 Gabriel, P., Marco, C., 2019. Computational optimal transport with
1183 applications to data sciences. Foundations and Trends® in Machine
1184 Learning 11, 355–607. URL: <http://dx.doi.org/10.1561/22000000073>,
1185 doi:10.1561/22000000073.
- 1186 Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D.,
1187 Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks.
1188 Communications of the ACM 63, 139–144. URL: [http://dx.doi.org/10.](http://dx.doi.org/10.1145/3422622)
1189 [1145/3422622](http://dx.doi.org/10.1145/3422622), doi:10.1145/3422622.
- 1190 Graffigna, M.S., Bartol, I.R., Tano, M.E., Dewji, S.A., 2026. Exploratory
1191 application of dmd for particle deposition and fluid field in the respira-
1192 tory tract. Journal of Aerosol Science 191, 106718. URL: [http://dx.](http://dx.doi.org/10.1016/j.jaerosci.2025.106718)
1193 [doi.org/10.1016/j.jaerosci.2025.106718](http://dx.doi.org/10.1016/j.jaerosci.2025.106718), doi:10.1016/j.jaerosci.
1194 [2025.106718](http://dx.doi.org/10.1016/j.jaerosci.2025.106718).
- 1195 Guo, X., Li, W., Iorio, F., 2016. Convolutional neural networks for steady
1196 flow approximation, in: Proceedings of the 22nd ACM SIGKDD Inter-
1197 national Conference on Knowledge Discovery and Data Mining, ACM. p.
1198 481–490. URL: <http://dx.doi.org/10.1145/2939672.2939738>, doi:10.
1199 [1145/2939672.2939738](http://dx.doi.org/10.1145/2939672.2939738).
- 1200 Han, M., Sane, S., Johnson, C.R., 2022. Exploratory lagrangian-
1201 based particle tracing using deep learning. Journal of Flow
1202 Visualization and Image Processing 29, 73–96. URL: [http:](http://dx.doi.org/10.1615/JFlowVisImageProc.2022041197)
1203 [//dx.doi.org/10.1615/JFlowVisImageProc.2022041197](http://dx.doi.org/10.1615/JFlowVisImageProc.2022041197), doi:10.1615/
1204 [jflowvisimageproc.2022041197](http://dx.doi.org/10.1615/JFlowVisImageProc.2022041197).
- 1205 Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models,
1206 in: Proceedings of the 34th International Conference on Neural Informa-
1207 tion Processing Systems, Curran Associates Inc., Red Hook, NY, USA.

1208 Hofmann, W., 2011. Modelling inhaled particle deposition in the hu-
1209 man lung—a review. *Journal of Aerosol Science* 42, 693–724. URL:
1210 <https://dx.doi.org/10.1016/j.jaerosci.2011.05.007>, doi:10.1016/
1211 j.jaerosci.2011.05.007.

1212 Huang, F., Zhu, Q., Zhou, X., Gou, D., Yu, J., Li, R., Tong, Z., Yang, R.,
1213 2021. Role of cfd based in silico modelling in establishing an in vitro-in vivo
1214 correlation of aerosol deposition in the respiratory tract. *Advanced Drug*
1215 *Delivery Reviews* 170, 369–385. URL: [http://dx.doi.org/10.1016/j.](http://dx.doi.org/10.1016/j.addr.2020.09.007)
1216 [addr.2020.09.007](http://dx.doi.org/10.1016/j.addr.2020.09.007), doi:10.1016/j.addr.2020.09.007.

1217 Huhn, Q.A., Tano, M.E., Ragusa, J.C., Choi, Y., 2023. Parametric dynamic
1218 mode decomposition for reduced order modeling. *Journal of Computational*
1219 *Physics* 475, 111852. doi:10.1016/j.jcp.2022.111852.

1220 ICRP, 1994. Human respiratory tract model for radiological protection icrp
1221 publication 66. Oxford: Pergamon .

1222 Kingma, D.P., Welling, M., 2014. Auto-encoding variational bayes, in:
1223 *International Conference on Learning Representations*. URL: [https:](https://openreview.net/forum?id=33X9fd2-9FyZd)
1224 [//openreview.net/forum?id=33X9fd2-9FyZd](https://openreview.net/forum?id=33X9fd2-9FyZd).

1225 Kleinstreuer, C., Zhang, Z., 2010. Airflow and particle transport in the hu-
1226 man respiratory system. *Annual Review of Fluid Mechanics* 42, 301–334.
1227 URL: <https://dx.doi.org/10.1146/annurev-fluid-121108-145453>,
1228 doi:10.1146/annurev-fluid-121108-145453.

1229 Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P., Hoyer, S.,
1230 2021. Machine learning–accelerated computational fluid dynamics. *Pro-*
1231 *ceedings of the National Academy of Sciences* 118. URL: [http://dx.doi.](http://dx.doi.org/10.1073/pnas.2101784118)
1232 [org/10.1073/pnas.2101784118](http://dx.doi.org/10.1073/pnas.2101784118), doi:10.1073/pnas.2101784118.

1233 Kundu, D., Panchagnula, M.V., 2025. Stochastic asymmetric bronchial tree
1234 models for population-scale variability in dosimetry. *Journal of Aerosol*
1235 *Science* 189, 106622. URL: [http://dx.doi.org/10.1016/j.jaerosci.](http://dx.doi.org/10.1016/j.jaerosci.2025.106622)
1236 [2025.106622](http://dx.doi.org/10.1016/j.jaerosci.2025.106622), doi:10.1016/j.jaerosci.2025.106622.

1237 Kutz, J.N., Brunton, S.L., 2016. *Dynamic Mode Decomposition*. Society for
1238 *Industrial and Applied Mathematics*. doi:10.1137/1.9781611974508.

- 1239 Kutz, J.N., Fu, X., Brunton, S.L., 2016. Multiresolution dynamic mode
1240 decomposition. *SIAM Journal on Applied Dynamical Systems* 15,
1241 713–735. URL: <http://dx.doi.org/10.1137/15M1023543>, doi:10.1137/
1242 15m1023543.
- 1243 Labiris, N.R., Dolovich, M.B., 2003. Pulmonary drug delivery.
1244 part i: Physiological factors affecting therapeutic effectiveness of
1245 aerosolized medications. *British Journal of Clinical Pharmacology*
1246 56, 588–599. URL: <http://dx.doi.org/10.1046/j.1365-2125.2003.01892.x>,
1247 doi:10.1046/j.1365-2125.2003.01892.x.
- 1248 Langtry, R.B., Menter, F.R., Likki, S.R., Suzen, Y.B., Huang, P.G., Völker,
1249 S., 2004. A correlation-based transition model using local variables—part
1250 ii: Test cases and industrial applications. *Journal of Turbomachinery* 128,
1251 423–434. URL: <http://dx.doi.org/10.1115/1.2184353>, doi:10.1115/
1252 1.2184353.
- 1253 Li, T., Biferale, L., Bonaccorso, F., Scarpolini, M.A., Buzzicotti, M., 2024a.
1254 Synthetic lagrangian turbulence by generative diffusion models. *Nature*
1255 *Machine Intelligence* 6, 393–403. URL: <http://dx.doi.org/10.1038/s42256-024-00810-0>,
1256 doi:10.1038/s42256-024-00810-0.
- 1257 Li, T., Biferale, L., Bonaccorso, F., Scarpolini, M.A., Buzzicotti, M.,
1258 2024b. Synthetic lagrangian turbulence by generative diffusion models.
1259 *Nature Machine Intelligence* URL: <https://www.nature.com/articles/s42256-024-00810-0>,
1260 doi:10.1038/s42256-024-00810-0.
- 1261 Li, Z., Kovachki, N.B., Azizzadenesheli, K., liu, B., Bhattacharya, K., Stuart,
1262 A., Anandkumar, A., 2021. Fourier neural operator for parametric partial
1263 differential equations, in: *International Conference on Learning Represen-*
1264 *tations*. URL: <https://openreview.net/forum?id=c8P9NQVtmn0>.
- 1265 Liu, H., Ma, S., Hu, T., Ma, D., 2023a. Computational investigation of
1266 flow characteristics and particle deposition patterns in a realistic human
1267 airway model under different breathing conditions. *Respiratory Physiology*
1268 *& Neurobiology* 314, 104085. URL: <https://doi.org/10.1016/j.resp.2023.104085>,
1269 doi:10.1016/j.resp.2023.104085.
- 1270 Liu, J., Anirudh, R., Thiagarajan, J.J., He, S., Mohan, K.A., Kamilov, U.S.,
1271 Kim, H., 2023b. Dolce: A model-based probabilistic diffusion framework

1272 for limited-angle ct reconstruction, in: Proceedings of the IEEE/CVF In-
1273 ternational Conference on Computer Vision (ICCV), pp. 10498–10508.

1274 Liu, K., Qiu, G., Tang, W., Zhou, F., 2020. Spectral regularization for
1275 combating mode collapse in gans. *Image and Vision Computing* 104,
1276 104005. URL: <http://dx.doi.org/10.1016/j.imavis.2020.104005>,
1277 doi:10.1016/j.imavis.2020.104005.

1278 Longest, P.W., Bass, K., Dutta, R., Rani, V., Thomas, M.L., El-Achwah,
1279 A., Hindle, M., 2019. Use of computational fluid dynamics deposition
1280 modeling in respiratory drug delivery. *Expert Opinion on Drug Delivery*
1281 16, 7–26. URL: <http://dx.doi.org/10.1080/17425247.2019.1551875>,
1282 doi:10.1080/17425247.2019.1551875.

1283 Longest, P.W., Holbrook, L.T., 2012. In silico models of aerosol delivery to
1284 the respiratory tract — development and applications. *Advanced Drug*
1285 *Delivery Reviews* 64, 296–311. URL: [https://dx.doi.org/10.1016/j.](https://dx.doi.org/10.1016/j.addr.2011.05.009)
1286 [addr.2011.05.009](https://dx.doi.org/10.1016/j.addr.2011.05.009), doi:10.1016/j.addr.2011.05.009.

1287 Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E., 2021. Learn-
1288 ing nonlinear operators via deeponet based on the universal approxi-
1289 mation theorem of operators. *Nature Machine Intelligence* 3, 218–229.
1290 URL: <http://dx.doi.org/10.1038/s42256-021-00302-5>, doi:10.1038/
1291 s42256-021-00302-5.

1292 Massey, F.J., 1951. The kolmogorov-smirnov test for goodness of fit. *Journal*
1293 *of the American Statistical Association* 46, 68–78. URL: [http://dx.doi.](http://dx.doi.org/10.1080/01621459.1951.10500769)
1294 [org/10.1080/01621459.1951.10500769](http://dx.doi.org/10.1080/01621459.1951.10500769), doi:10.1080/01621459.1951.
1295 10500769.

1296 Menter, F.R., 1994. Two-equation eddy-viscosity turbulence models for
1297 engineering applications. *AIAA Journal* 32, 1598–1605. URL: [http:](http://dx.doi.org/10.2514/3.12149)
1298 [//dx.doi.org/10.2514/3.12149](http://dx.doi.org/10.2514/3.12149), doi:10.2514/3.12149.

1299 Menter, F.R., Langtry, R., Völker, S., 2006a. Transition modelling for general
1300 purpose cfd codes. *Flow, turbulence and combustion* 77, 277–303.

1301 Menter, F.R., Langtry, R.B., Likki, S.R., Suzen, Y.B., Huang, P.G., Volker,
1302 S., 2006b. A correlation-based transition model using local variables—part
1303 i: Model formulation. *Journal of Turbomachinery* 128, 413. URL: [http:](http://dx.doi.org/10.1115/1.2184352)
1304 [//dx.doi.org/10.1115/1.2184352](http://dx.doi.org/10.1115/1.2184352), doi:10.1115/1.2184352.

- 1305 Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A., 2019.
 1306 Occupancy networks: Learning 3d reconstruction in function space, in:
 1307 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition
 1308 (CVPR), IEEE. p. 4455–4465. URL: [http://dx.doi.org/10.1109/CVPR.](http://dx.doi.org/10.1109/CVPR.2019.00459)
 1309 2019.00459, doi:10.1109/cvpr.2019.00459.
- 1310 Nichol, A.Q., Dhariwal, P., 2021. Improved denoising diffusion probabilistic
 1311 models, in: Meila, M., Zhang, T. (Eds.), Proceedings of the 38th Interna-
 1312 tional Conference on Machine Learning, PMLR. pp. 8162–8171.
- 1313 Nowak, N., Kakade, P.P., Annapragada, A.V., 2003. Computational fluid
 1314 dynamics simulation of airflow and aerosol deposition in human lungs.
 1315 Annals of Biomedical Engineering 31, 374–390. URL: [http://dx.doi.](http://dx.doi.org/10.1114/1.1560632)
 1316 [org/10.1114/1.1560632](http://dx.doi.org/10.1114/1.1560632), doi:10.1114/1.1560632.
- 1317 Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S., 2019.
 1318 DeepSDF: Learning continuous signed distance functions for shape repre-
 1319 sentation, in: 2019 IEEE/CVF Conference on Computer Vision and Pat-
 1320 tern Recognition (CVPR), IEEE. p. 165–174. URL: [http://dx.doi.org/](http://dx.doi.org/10.1109/CVPR.2019.00025)
 1321 [10.1109/CVPR.2019.00025](http://dx.doi.org/10.1109/CVPR.2019.00025), doi:10.1109/cvpr.2019.00025.
- 1322 Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A., 2018.
 1323 Film: visual reasoning with a general conditioning layer, in: Proceed-
 1324 ings of the Thirty-Second AAAI Conference on Artificial Intelligence and
 1325 Thirtieth Innovative Applications of Artificial Intelligence Conference and
 1326 Eighth AAAI Symposium on Educational Advances in Artificial Intelli-
 1327 gence, AAAI Press.
- 1328 Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P., 2021. Learn-
 1329 ing mesh-based simulation with graph networks, in: International Con-
 1330 ference on Learning Representations. URL: [https://openreview.net/](https://openreview.net/forum?id=roNqYLO_XP)
 1331 [forum?id=roNqYLO_XP](https://openreview.net/forum?id=roNqYLO_XP).
- 1332 Poorbahrami, K., Mummy, D.G., Fain, S.B., Oakes, J.M., 2019. Patient-
 1333 specific modeling of aerosol delivery in healthy and asthmatic adults.
 1334 Journal of Applied Physiology 127, 1720–1732. URL: [http://dx.doi.](http://dx.doi.org/10.1152/jappphysiol.00221.2019)
 1335 [org/10.1152/jappphysiol.00221.2019](http://dx.doi.org/10.1152/jappphysiol.00221.2019), doi:10.1152/jappphysiol.
 1336 00221.2019.

- 1337 Ramdas, A., Trillos, N., Cuturi, M., 2017. On wasserstein two-sample testing
1338 and related families of nonparametric tests. *Entropy* 19, 47. URL: <http://dx.doi.org/10.3390/e19020047>, doi:10.3390/e19020047.
1339
- 1340 Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B., 2022.
1341 High-resolution image synthesis with latent diffusion models, in: 2022
1342 IEEE/CVF Conference on Computer Vision and Pattern Recognition
1343 (CVPR), IEEE. p. 10674–10685. URL: <http://dx.doi.org/10.1109/CVPR52688.2022.01042>, doi:10.1109/cvpr52688.2022.01042.
1344
- 1345 Rostami, A.A., 2009. Computational modeling of aerosol deposition
1346 in respiratory tract: A review. *Inhalation Toxicology* 21, 262–290.
1347 URL: <http://dx.doi.org/10.1080/08958370802448987>, doi:10.1080/
1348 08958370802448987.
- 1349 ROWLEY, C.W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P., HEN-
1350 NINGSON, D.S., 2009. Spectral analysis of nonlinear flows. *Journal*
1351 *of Fluid Mechanics* 641, 115–127. URL: [http://dx.doi.org/10.1017/](http://dx.doi.org/10.1017/S0022112009992059)
1352 [S0022112009992059](http://dx.doi.org/10.1017/S0022112009992059), doi:10.1017/s0022112009992059.
- 1353 SCHMID, P.J., 2010. Dynamic mode decomposition of numerical and experi-
1354 mental data. *Journal of Fluid Mechanics* 656, 5–28. URL: [http://dx.doi.](http://dx.doi.org/10.1017/S0022112010001217)
1355 [org/10.1017/S0022112010001217](http://dx.doi.org/10.1017/S0022112010001217), doi:10.1017/s0022112010001217.
- 1356 Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S., 2015. Deep
1357 unsupervised learning using nonequilibrium thermodynamics, in: Bach,
1358 F., Blei, D. (Eds.), *Proceedings of the 32nd International Conference on*
1359 *Machine Learning*, PMLR, Lille, France. pp. 2256–2265. URL: [https:](https://proceedings.mlr.press/v37/sohl-dickstein15.html)
1360 [//proceedings.mlr.press/v37/sohl-dickstein15.html](https://proceedings.mlr.press/v37/sohl-dickstein15.html).
- 1361 Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole,
1362 B., 2021. Score-based generative modeling through stochastic differen-
1363 tial equations, in: *International Conference on Learning Representations*
1364 (ICLR). URL: <https://arxiv.org/abs/2011.13456>.
- 1365 Székely, G.J., Rizzo, M.L., 2013. Energy statistics: A class of statistics based
1366 on distances. *Journal of Statistical Planning and Inference* 143, 1249–1272.
1367 URL: <http://dx.doi.org/10.1016/j.jspi.2013.03.018>, doi:10.1016/
1368 [j.jspi.2013.03.018](http://dx.doi.org/10.1016/j.jspi.2013.03.018).

- 1369 Talaat, K., Xi, J., Baldez, P., Hecht, A., 2019. Radiation dosimetry of inhaled
1370 radioactive aerosols: Cfpd and mcnp transport simulations of radionuclides
1371 in the lung. Scientific Reports 9. URL: [https://dx.doi.org/10.1038/
1372 s41598-019-54040-1](https://dx.doi.org/10.1038/s41598-019-54040-1), doi:10.1038/s41598-019-54040-1.
- 1373 Zbilut, J.P., Marwan, N., 2008. The wiener–khinchin theorem and recur-
1374 rence quantification. Physics Letters A 372, 6622–6626. URL: [http://dx.
1375 doi.org/10.1016/j.physleta.2008.09.027](http://dx.doi.org/10.1016/j.physleta.2008.09.027), doi:10.1016/j.physleta.
1376 2008.09.027.