

Validated respiratory drug deposition predictions from 2D and 3D medical images with statistical shape models and convolutional neural networks

Josh Williams¹, Haavard Ahlqvist¹, Alexander Cunningham¹, Andrew Kirby², Ira Katz³, John Fleming^{4,5}, Joy Conway^{4,6}, Steve Cunningham⁷, Ali Ozel^{1*}, and Uwe Wolfram^{1*}

¹School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK

²Royal Hospital for Children and Young People, NHS Lothian, Edinburgh, UK

³Consultant, Meudon, France

⁴National Institute of Health Research Biomedical Research Centre in Respiratory Disease, Southampton, UK

⁵Department of Medical Physics and Bioengineering, University Hospital Southampton NHS Foundation Trust, Southampton, UK

⁶Respiratory Sciences, Centre for Health and Life Sciences, Brunel University, London, UK

⁷Centre for Inflammation Research, University of Edinburgh, Edinburgh, UK

March 3, 2023

* UW and AO share last authorship.

Address correspondence to u.wolfram@hw.ac.uk or a.ozel@hw.ac.uk

Abstract

For the one billion sufferers of respiratory disease, managing their disease with inhalers crucially influences their quality of life. Generic treatment plans could be improved with the aid of computational models that account for patient-specific features such as breathing pattern, lung pathology and morphology. Therefore, we aim to develop and validate an automated computational framework for patient-specific deposition modelling. To that end, an image processing approach is proposed that could produce 3D patient respiratory geometries from 2D chest X-rays and 3D CT images. We evaluated the airway and lung morphology produced by our image processing framework, and assessed deposition compared to *in vivo* data. The 2D-to-3D image processing reproduces airway diameter to 9% median error compared to ground truth segmentations, but is sensitive to outliers of up to 33% due to lung outline noise. Predicted regional deposition gave 5% median error compared to *in vivo* measurements. The proposed framework is capable of providing patient-specific deposition measurements for varying treatments, to determine which treatment would best satisfy the needs imposed by each patient (such as disease and lung/airway morphology). Integration of patient-specific modelling into clinical practice as an additional decision-making tool could optimise treatment plans and lower the burden of respiratory diseases.

1 Introduction

Respiratory diseases affect over one billion people worldwide (The Global Asthma Network, 2018). Diseases such as asthma, cystic fibrosis and COPD are typically treated with inhaled therapeutics which are delivered to the affected tissue to prevent symptoms such as breathlessness from arising, or relieve them when they do arrive. Many patients' symptoms are worsened by a combination of poor inhaler technique and a lack of adherence to their treatment plan (van Boven et al., 2018). Electronic health (e-health) measures such as smart inhalers (van Boven et al., 2018) are being introduced to improve adherence and quality of care (Honkoop et al., 2022). Smart inhalers can use sensors to track how often patients use inhalers and also provide estimates on technique or inspiratory flowrate based on acoustic sensors (van Boven et al., 2018; Lim et al., 2019). The value of e-health monitoring for patients could be increased significantly by linking measurements such as inhalation flowrate with patient-specific physical models that accurately predict how the inhaled drug has deposited in patient lungs.

Patient-specific modelling of drug deposition can be done with computational particle-fluid dynamics (CPFD) which solves the governing equations of air and drug particle transport during inhalation (Nowak et al., 2003; Koullapis et al., 2016; Williams et al., 2022a). Patient-specific CPFD modelling requires 3D representations of the patient's respiratory system, usually obtained from clinical computed tomography (CT) scans (Kleinstreuer and Zhang, 2010). Typically the airways are segmented using manual approaches, which can take several hours (Sonka et al., 1996), and is therefore unsuitable for usage in future e-health monitoring frameworks. Tschirren et al. (2009) showed that manually correcting airway segmentations obtained from semi-automatic methods can also take several hours. Such a long and labour-intensive process is a substantial bottleneck in applying patient-specific deposition models to a large number of respiratory patients. Fully-automatic segmentation methods have recently become available, which are promising for rapid segmentation of lungs and airways from volumetric CT using convolutional neural networks (CNNs) (Garcia-Uceda et al., 2021; Hofmanninger et al., 2020; Yun et al., 2019; Tang et al., 2019). CNNs apply varying levels of convolution, where the convolution kernels are learned by evaluating the difference between the CNN output compared to a 'ground truth' human-verified

result (LeCun et al., 1995). CNNs have been widely used in medical image processing and are used here to segment airways from 3D images such as clinical CT. However, CT images are more costly to acquire than chest X-rays and are therefore not optimal for imaging a large number of respiratory sufferers. Additionally, using CT scans is not encouraged in children for routine assessments due to harmful radiation. Therefore, an additional image-processing approach that can extract 3D respiratory systems from sparser imaging such as chest X-rays is an unmet need.

To perform deposition simulations based on patient chest X-rays, we require a novel image-processing approach to reconstruct 3D respiratory geometries from 2D chest X-rays. Reconstruction of 3D skeletal anatomies including femora, vertebra or knee joints from planar and bi-planar X-ray images has been performed using statistical shape and appearance models (SSAMs) (Albrecht et al., 2013; Baka et al., 2011; Väänänen et al., 2015; Zheng et al., 2011; Youn et al., 2017; Dworzak et al., 2010). Statistical shape models (SSMs) take a set of corresponding geometries (usually represented with landmarks that represent key features), and quantifies how shape features change across a population (based on eigenvalues and eigenvectors of the covariance matrix) (Cootes et al., 1995). A SSAM performs the same operation, including an additional feature such as gray value at each landmark. In this approach, patient lungs may be extracted from a chest X-ray by combining (i) available information such as lung outline and gray-value distribution, as well as (ii) a SSAM that describes how lung shapes and projected gray-value varies from a database of known (3D) lung shapes (Cootes et al., 1995, 2001; Heimann and Meinzer, 2009). Respiratory statistical shape models (SSMs) have previously been used for diagnostic purposes, as Irving et al. (2013) developed an approach to extract 3D airways from a single chest X-ray and diagnose tuberculosis in children (Irving et al., 2011, 2013). Osanlouy et al. (2020) used a SSM of the lungs to quantify differences in lung morphology in smokers and non-smokers. These examples of SSM and SSAM applications in skeletal and respiratory imaging show that SSAMs are a suitable approach for extracting full respiratory systems from planar or bi-planar chest X-rays. However, a key issue in this approach is that only the first bifurcation-level (or ‘generation’) of the airway tree is visible on a chest X-ray (Irving et al., 2013), as the low density lung structures are overlapped by denser structures such as the spine, rib cage and heart

(Gozes and Greenspan, 2018; Yang et al., 2017; Liang et al., 2020). Similarly, only around 5-10 generations (out of 23) are visible on high resolution CT scans (Kleinstreuer and Zhang, 2010; Bordas et al., 2015). Thus, even with high resolution imaging, the deep lung airways are left unaccounted for and must be included for patient-specific deposition predictions.

The limitations of imaging resolution have been overcome in several ways. The remaining levels of the conducting airway tree can be generated (approximately up to generation 16) using statistical models based on information from the visible airways (Tawhai et al., 2004; Bordas et al., 2015; Montesantos et al., 2016). The remaining portion of the airway tree (bifurcations 17 to 23) are the respiratory ‘acinar’ region, which is a very small (characteristic length is $\mathcal{O}(100\ \mu\text{m})$) and intricate structure that is impossible to image *in vivo* (Haefeli-Bleuer and Weibel, 1988; Rausch et al., 2011; Dong et al., 2022). Acinar flows have been studied with CPFDF (Sznitman, 2013; Dong et al., 2022; Khajeh-Hosseini-Dalasm and Longest, 2015; Hofemeier and Sznitman, 2016), but it is highly computationally expensive to directly resolve the full conducting airways and deep lung in one simulation (Koullapis et al., 2018a, 2019, 2020). Typically, the deep lung mechanics are treated with reduced-dimension models such as 1D airway networks or 0D lumped parameter models which are coupled to 3D CPFDF simulations at the 3D airway outlet (Oakes et al., 2015, 2018; Roth et al., 2017; Comerford et al., 2010; Ismail et al., 2014; Pozin et al., 2017; Kuprat et al., 2020; Poorbahrami and Oakes, 2019; Feng et al., 2021). To model the full respiratory system, we therefore couple 3D CPFDF domains produced from our image-processing framework with a 0D lumped parameter model to validate the entire image-processing and simulation framework against experimental deposition data (Conway et al., 2012; Fleming et al., 2015).

Therefore, we aimed to (i) develop a computational framework to reconstruct 3D airways and lungs from 2D image(s); (ii) develop a segmentation tool for reconstructing airways and lungs from clinical CT data; (iii) evaluate the ability of the developed image-processing tools to reproduce morphological properties and deposition of the ground truth airways, and (iv) validate deposition predictions against experimental data. Our developed computational framework is made publicly available online (<https://github.com/jvwilliams23/respiratory2Dto3Dpaper>).

2 Methods and materials

2.1 Patient image data and ground truth generation

We used 51 clinical CT scans from the LUNA16 challenge (Setio et al., 2017), which is part of The Cancer Imaging Archive (TCIA) (Clark et al., 2013). The scans were acquired for lung cancer nodule detection which does not depend on inhalation state, therefore no information on breath-hold procedures was reported by Setio et al. (2017). The scans used were chosen randomly by Tang et al. (2019), and they provided radiologist-verified lung lobe segmentations online. We created airway segmentations for the same scans using a semi-automatic region-growing approach based on Nardelli et al. (2015). These segmentations have been independently verified by a radiologist with 9 years experience. We excluded 13 samples from our SSAM dataset, as the segmentations only included the trachea and first bifurcation.

To validate our modelling approach, we compared to experimental data of radiopharmaceutical aerosol deposition from Conway et al. (2012) and Fleming et al. (2015), which we refer to throughout as the ‘Southampton/Air Liquide dataset’. Conway et al. (2012) and Fleming et al. (2015) performed experiments with inhaled aerosol in six healthy and six asthmatic subjects. Patients sat in an erect (sitting) position while aerosol was delivered through an AKITA nebuliser. The AKITA nebuliser is breath-actuated to release drug only during inhalation, and a controlled (constant) inhalation profile to optimise deposition (Munro et al., 2020). Deposition was imaged with single photon emission computed tomography (SPECT) and combined with low resolution X-ray CT captured in the same machine (SPECT-CT), which relates deposition images to anatomical structure of the lungs (Fleming et al., 2011). To relate deposition measurements to airway morphology, high resolution CT was acquired for each patient on a different machine on a different day (Conway et al., 2012; Montesantos et al., 2013). Full information on scanning equipment and protocol are provided elsewhere (Conway et al., 2012; Montesantos et al., 2013), and only key points related to our study are discussed here. High resolution CT images were acquired in supine position, which is known to decrease lung volume compared to erect position (Ibanez and Raurich, 1982). Additionally, Conway et al. (2012) discussed imaging issues in the upper airway. In half of the images acquired, the epiglottis was completely closed, meaning there was no path for air to

pass from the mouth to the lung airways. This may have occurred since the images were taken during exhalation, where the glottis area is at its narrowest (Brancatisano et al., 1983; Scheinherr et al., 2015). Furthermore, motion artefacts may have blurred the already narrow glottis to appear fully closed. To allow for air and drug particles to pass from the mouth to the lung, we manually cleaned up segmentations in these problematic cases (Supplementary Material Section 1).

2.2 Shape reconstruction from a chest X-ray image

To reconstruct 3D lung and airway geometries from an X-ray image, we created a statistical shape and appearance model (SSAM) based on ground truth airway and lung lobe segmentations. The appearance is based on digitally reconstructed radiographs (DRRs) derived from the patient CT data (described below). The DRRs were also used for testing how well the SSAM fits to an X-ray. Below we describe the process to create a SSAM from patient CT data and segmentations described above (Section 2.1).

2.2.1 Pre-processing and landmarking

To establish correspondence amongst the lobes, we make use of the '*Growing and Adaptive Meshes*' (GAMEs) algorithm of Ferrarini et al. (2007). This algorithm determines the set of landmarks needed to best describe the surface mesh by clustering points based on their proximity and covariance. A set of landmarks are grown for a base mesh that acts as a population mean (Marsland et al., 2002). The base landmarks are then iteratively adapted to fit the rest of the population, such that adapted points correspond to the same feature (Kohonen, 1990, 2013). Outlier points that had not correctly fitted to a new shape were detected as points with a Mahalanobis distance of greater than 0.15, which we determined empirically by analysing statistics of poorly fitted cases. Once all points were below this threshold, the landmark adaptation was considered converged. If convergence was not satisfied and the maximum distance has settled at a steady-state (change in maximum Mahalanobis distance less than 0.05 after five full iterations through all points in ground truth), then the landmarking was considered failed. If landmarking failed, another cloud of landmark points was used as a base to adapt from.

As we cannot directly see the lobar fissures on a chest X-ray, we removed landmarks on the lobar fissures from the landmark list. Landmarks were classified as being on a lobar fissure if the (dimensionless) Mahalanobis distance between a landmark in one lobe and a landmark in the adjacent lobe was less than 0.2. This allowed the SSAM to capture more variance on the outer curvature of the lungs instead of the highly variable lobar fissures, which improves adaptation to lung outlines on X-ray images. Landmarks for all lobes were combined into one dataset to have a SSAM for the entire lungs (instead of individual lobes), which allows us to capture inter-organ relations (Cerrolaza et al., 2019) such as distance between both lungs, orientation and non-overlapping structure of lobes. This gave us 2966 landmarks in total for the lungs.

To landmark the airways we first skeletonised the segmentation in scikit-image (Lee et al., 1994). The skeletonised labelmap was then converted to spatial coordinates with the Skan package (Palmer et al., 2013). A graph was created from coordinates output from Skan using Networkx (Hagberg et al., 2008). Branches with length less than 5 mm were removed as recommended by Tschirren et al. (2005b). We included the trachea, main and lobar bronchi in the skeleton landmarks, as airways below this will not be visible on the X-ray image.

At each node on the graph we computed the diameter based on the cross sectional area of the 3D segmentation at that location. This was used to get landmarks on the surface by creating a circle centered on and perpendicular to the medial axis. We computed a vector from the circle center to each point along its circumference. We found where this vector intersected the surface mesh of the segmentation and set a landmark here. This allowed us to landmark irregular cross-sections. The airway landmarks were combined with lung landmarks, giving us 3486 landmarks in total.

We created DRRs for all patients in the LUNA16 dataset which we used to train the SSAM and assessing performance. This was done by integrating the Hounsfield units for each voxel in the anterior-posterior direction of the scan to create a frontal chest X-ray. We also created additional DRRs in the sagittal plane to test the influence of including an additional projection on the morphologies predicted by the SSAM. The gray-value at the landmark location could then be used in the appearance part of the SSAM. The DRRs represent

patient X-rays at the same time point that the CT data used to extract landmarks was taken. The DRRs were also used for assessing our SSAM performance, as we can compare the 3D reconstructions of the lungs and airways from the SSAM to CT at the same instant. By comparing lungs and airways at the same time instant, our results are not influenced by changes in lung or airway shape during breathing.

Similar to Baka et al. (2011), we computed a Canny edge map of the chest X-ray to use for guiding our SSAM fitting. To reduce noise in the edge map, we first downsampled the X-ray to be four times coarser and applied global histogram equalization. We then applied a Canny edge map (Canny, 1986), where the Gaussian kernel has a width of two pixels. We found this combination of pre-processing steps to produce an edge map with the lowest amount of noise from overlapping structures such as the rib cage.

2.2.2 Statistical shape and appearance modelling

Once a set of landmarks and DRRs were extracted from the LUNA16 dataset (Section 2.2.1), all sets of landmarks were aligned to zero mean and isotropic scaling was applied such that the coordinates had zero mean and unit standard deviation. Landmark gray values were also normalised to zero mean and unit standard deviation per sample. We then create a SSAM to describe and parameterise the covariance of the lung and airway shape-appearance data in the LUNA16 dataset. The SSAM parameters are then iteratively adjusted based on how well the lung and airway shape and appearance agrees with the X-ray image (Figure 1a-b).

The SSAM uses principal component analysis (PCA) to create a model of the variation in the shape and the gray-values of the corresponding chest X-ray. The shape and gray-value, x can then be described by the linear function

$$x \approx \bar{x} + \sum_{m=1}^{N_m} \Phi_m \cdot \left(b_m \sqrt{\sigma_m^2} \right), \quad (1)$$

where \bar{x} is a vector of mean shape and gray-value. It has a size of $(3 + N_{XR})N_L$, where N_L is number of landmarks, corresponding to a Cartesian coordinate for each landmark and N_{XR} is the number of X-rays per patient, which each contains a set of gray-values for each landmark. Φ is a $(3 + N_{XR})N_L \times N_m$ matrix, giving the main modes of variation of the

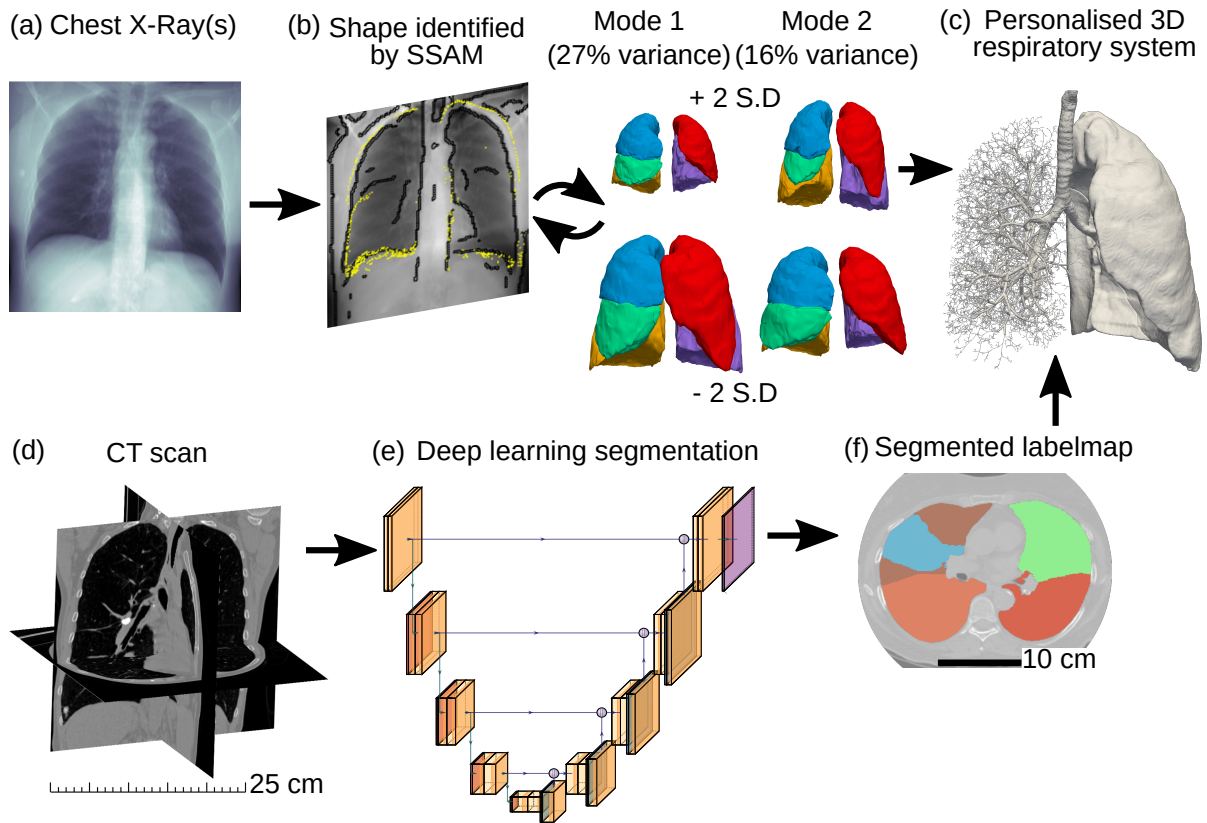


Figure 1: Schematic overview of image-processing framework. Chest X-ray image(s) are provided as an input (a), where the shape is identified by our SSAM by iteratively adapting lung outline and gray-value based on a database of known lung shapes (b-c). The SSAM reconstruction can then be used to generate a surface mesh of lungs with a full airway tree (c). Alternatively, a 3D CT scan may be provided as an input (d) to our trained CNN architecture (e). The CNN returns a segmented labelmap for lungs and airways (f), which can also be used to generate the full respiratory domain (c).




Figure 2: Assessment of sensitivity of SSAM metrics to number of modes included in the model. Panels show (a) explained variance based on number of modes, (b) reconstruction error for shapes in training set, (c) generalisation error for unseen shapes. When the model describes 90% of the population variance (indicated by blue dashed line), the generalisation error in adapting the model to unseen shapes is saturated.

training set, where N_m is the number of modes. σ_m^2 is the variance described by mode m . b_m is the shape parameter for mode m , where Figure 1(b) shows the effect of varying the weighting of b for the first two modes. By using PCA, the complexity of the model can be reduced by including only the number of modes, N_m , that are required to capture a specified amount of variance in the modelled shape. In this study, 90% of the population variance is included in the model ($N_m = 25$, Figure 2a). The SSAM was created using our Python library ‘pyssam’ for statistical shape and appearance modelling (Williams et al., 2023).

The accuracy obtained by the SSAM with N_m modes can be assessed in terms of its ability to reconstruct a shape in the training dataset (reconstruction error). More importantly, the SSAM can be assessed by its ability to generalise and reconstruct an unseen shape (not in the training set), which is known as generalisation error. When the shape model contains 100% of the training population variance (Figure 2a), the reconstruction error should tend to zero as all information required to reconstruct the shape is contained within the model (Figure 2b). We assessed reconstruction and generalisation error with leave-one-out testing, which was calculated over the entire training set and repeated 30 times. The mean absolute error is presented, which is expressed as a percentage of the lung bounding-box size. We observed that $N_m = 25$ is a suitable choice to reduce the SSAM dimensionality, as the generalisation error did not increase further by including additional modes (Figure 2c).

The shape parameters, \mathbf{b} , can then be optimised with respect to the fit of the SSAM outline (\mathcal{L}_{fit}) and the similarity to the training set (\mathcal{L}_{prior}). The appearance model optimises the shape by minimising the gray-value loss (\mathcal{L}_g). Gray-value loss is calculated as the difference between the image gray-value at the landmark location and the expected gray-value from Equation 1 (Väänänen et al., 2015; Sarkalkan et al., 2014; Castro-Mateos et al., 2014). We also used the anatomical shadow loss (\mathcal{L}_{AS}) for airways as proposed by Irving et al. (2013). Anatomical shadow compares local gray-values inside a small control area inside and outside of the projected airway surface, which can enhance airway contrast compared to dense overlapping structures such as the spinal column. Therefore, the optimal isotropic scaling, translation and shape parameters are found by minimising the equation

$$\mathcal{L} = C_{fit}\mathcal{L}_{fit} + C_{prior}\mathcal{L}_{prior} + C_g\mathcal{L}_g + C_{AS}\mathcal{L}_{AS} \quad (2)$$

where C is a coefficient that controls the weight of each fitting term.

The fit of the projected landmarks outline (‘silhouette’ landmarks) to the lung outline on the X-ray image is found by

$$\mathcal{L}_{fit} = 1 - \frac{1}{N_{SL}} \sum_{i=1}^{N_{SL}} D_i \quad (3)$$

where N_{SL} is the number of silhouette landmarks, D_i is a function based on the distance between each silhouette landmark and the nearest X-ray image outline point. Silhouette landmarks were defined as landmarks where the nearest surface mesh point is shared by two faces, with one face normal pointing towards and another away from the image source. For example, for an anterior-posterior (frontal) chest X-ray image, the projection is normal to the y -direction and therefore the two faces nearest to the silhouette landmark will have positive and negative y components in the normal vector. The silhouette landmarks represent the outline of the shape normal to the projection (Baka et al., 2011).

Baka et al. (2011) normalised the outline distance to bounds of [0,1] by

$$D_i = \exp\left(-\frac{\min(\|\mathbf{x}_{SSM,i} - \mathbf{x}_{XR}\|)}{C}\right). \quad (4)$$

where $C = 5$ was used by Baka et al. (2011) and also in this study. \mathcal{L}_{fit} is not evaluated for the airways as the airway outline extracted from a Canny edge map is weak and unreliable

due to overlapping denser regions on the X-ray image.

Improbable shapes (relative to our training set) are penalised by the Mahalanobis distance between the created shape and the mean shape (Baka et al., 2011; Ferrarini et al., 2007) calculated by

$$\mathcal{L}_{prior} = \frac{1}{N_L} \sum_{i=1}^{N_L} \sqrt{(\mathbf{x}_i - \bar{\mathbf{x}}_i)^T \text{cov}[\mathbf{x}, \bar{\mathbf{x}}]^{-1} (\mathbf{x}_i - \bar{\mathbf{x}}_i)}, \quad (5)$$

where N_L is the total number of landmarks, $\text{cov}[\mathbf{x}, \bar{\mathbf{x}}]^{-1}$ is the inverse covariance matrix between \mathbf{x}_i and $\bar{\mathbf{x}}_i$.

We use the appearance model to minimise the difference between the modelled gray-value (g_{model}) and the gray-value at the landmark location (g_{target}). This is found by

$$\mathcal{L}_g = \frac{1}{N_L} \sum_{i=1}^{N_L} |g_{model,i} - g_{target,i}|. \quad (6)$$

Finally, the anatomical shadow loss is found for each silhouette landmark as proposed by Irving et al. (2013). This is done by creating two regions-of-interest (ROIs) at some distance $\pm s_{AS}$ normal to the border of the silhouette landmarks. Each ROI is a circle with radius r_{AS} , and the gray-value ‘inside’ and ‘outside’ of the airway projection are compared as

$$\mathcal{L}_{AS} = \frac{1}{N_{SL}} \sum_{i=1}^{N_{SL}} \frac{g_{inside,i} - g_{outside,i}}{g_{outside,i}}. \quad (7)$$

We found the optimal parameters to be $s_{AS} = 20$ pixels and $r_{AS} = 14$ pixels.

Shape parameters were optimised using the NGO (Nevergrad optimiser) algorithm (Meunier et al., 2021) in the Nevergrad Python library (Rapin and Teytaud, 2018). This is a competence map method, which adjusts the optimisation algorithm based on the number of optimisation iterations and number of parameters optimised. We chose this method as it can therefore adjust to more modes of variation as our dataset grows in future studies. We found this gradient-free optimisation to show good robustness against local minima. The shape parameters for each mode were initialised to $\mathbf{b}_m = 0$, and were bound to $\mathbf{b}_m \pm 3$.

SSAM hyperparameters in the loss function (Equation 2) which produced the lowest maximum error in reconstructed lung space volume were found through Gaussian process regression (Perdikaris and Karniadakis, 2016; Angelikopoulos et al., 2012), also known as

Kriging (Krige, 1951). We performed 100 optimisation iterations, using a subset of 11 randomly selected samples from the LUNA16 dataset. This gave the optimal model parameters to be $C_{fit} = 0.795$, $C_g = 0.687$ and $C_{prior} = 4.4 \times 10^{-4}$. The anatomical shadow loss coefficient was later found through a simple grid search as $C_{AS} = 0.2$.

New landmarks obtained from the SSAM were converted to a surface by morphing a template mesh with template landmarks to the new set of landmarks (Grassi et al., 2011). New mesh points ($\mathbf{p}_{i,new}$) are created from the new landmarks ($\mathbf{x}_{i,new}$), template mesh points ($\mathbf{p}_{i,template}$) and template landmarks ($\mathbf{x}_{j,template}$) by

$$\mathbf{p}_{i,new} = \mathbf{p}_{i,template} + \sum_j^{N_L} k(\mathbf{p}_{i,template}, \mathbf{x}_{j,template}) w_j \quad (8)$$

where k is a kernel used for morphing the mesh and w_j represents the weights controlling how much each landmark point is morphed. The weights are computed by solving (8) for w_j and setting $\mathbf{p}_{i,new} = \mathbf{x}_{j,new}$. We adjusted the surface morphing algorithm of Grassi et al. (2011) to use a Gaussian kernel, as we found it to produce a smoother deformation when morphing. The kernel is $k(a, b) = \exp(-||a - b||^2 / 2\sigma^2)$, where $\sigma = 0.3$ (Carr et al., 1997).

2.3 Segmentation from volumetric CT

To segment lung lobes and airways from volumetric CT scans (Figure 1d), we used convolutional neural networks (LeCun et al., 1995) such as the U-Net shown in Figure 1e (Ronneberger et al., 2015; Çiçek et al., 2016). We implemented the U-Net in PyTorch (Paszke et al., 2019). The U-Net architecture was detailed by Ronneberger et al. (2015). For lung segmentation we used the pre-trained U-Net of Hofmanninger et al. (2020). For segmenting the airways we trained a U-Net using segmentations from the LUNA16 challenge (Figure 1e). Prior to training, all images and labelmaps were cropped to the extent of the airway labelmap to conserve memory and allow us to segment the entire 3D image, instead of slice-by-slice (Hofmanninger et al., 2020), or sliding-box approaches (Juarez et al., 2019). With the aim of lowering memory consumption, we also trained an ENet ('efficient neural network') as proposed by Paszke et al. (2016). The ENet has approximately 15 times fewer parameters than the U-Net architecture (Comelli et al., 2021), which creates lower memory consumption

during training and faster processing time for 3D images (Paszke et al., 2016; Lieman-Sifry et al., 2017).

The loss function used for training all CNNs was (Lin et al., 2017; Tang et al., 2019), given as

$$\mathcal{L} = \mathcal{L}_{DICE} + \mathcal{L}_{focal}, \quad (9)$$

where \mathcal{L}_{DICE} is the DICE loss and \mathcal{L}_{focal} is the focal loss. The DICE loss maximises the total intersection between a modelled labelmap and the ground-truth labelmap during training, and is found by

$$\mathcal{L}_{DICE} = \frac{1}{CN} \sum_c^C \sum_i^N \left(1 - \frac{p_{ic}g_{ic}}{p_{ic}g_{ic} + (1 - p_{ic})g_{ic} + p_{ic}(1 - g_{ic})} \right), \quad (10)$$

where p_{ic} is the probability of a voxel i belonging to class c as predicted by the CNN. g_{ic} is the ground truth value for the same voxel (which will be 1 or 0). This is averaged over all voxels in the image N and the total number of classes C .

The focal loss weights the loss more towards ‘harder to learn’ voxels with a focusing parameter, γ . Focal loss was developed for images with high class imbalances, such as airway segmentations, where we found the number of voxels representing the airways to be on average 63 times fewer than the number of background voxels. The focal loss was found by

$$\mathcal{L}_{focal} = -\frac{1}{CN} \sum_c^C \sum_i^N \alpha g_{ic} (1 - p_{ic})^\gamma \log(p_{ic}), \quad (11)$$

where we set the focusing parameter $\gamma = 5$ based on a parametric study of training with γ set to 1, 3 and 5 (Supplementary Material Figure S2). We used $\alpha = 1$ as recommended by Tang et al. (2019).

To train the U-Net and ENet, we used the Adam optimiser with a batch size equal to one to minimise memory consumption. Data was augmented with rotations of $\pm 15^\circ$. We created 50 new augmented datasets per epoch. Training was stopped once the validation loss had not decreased further in the most recent 10 epochs. To train the U-Net and ENet models, we used a compute node with 756 GB of RAM on Heriot-Watt University high performance computing cluster ROCKS. Training took 3 days for the U-Net (42 epochs) and 5 hours for the ENet (45 epochs). The ENet DICE coefficient was lower than the U-Net (0.89 compared

to 0.93, Figure S3a). The ENet had a significantly improved inference time compared to the U-Net (ENet gave 5.7 times speedup, Figure S3b).

2.4 Distal airway generation

To create the rest of the conducting airways not visible in CT or from our airway SSAM, we implemented a volume-filling airway generation algorithm (Bordas et al., 2015; Tawhai et al., 2004). The algorithm generates a full conducting airway tree based on information from the upper airways such as branch length and diameter. The lung space volume was discretised into seed points with a uniform spacing found by $\Delta = (V_{lung}/N_T)^{1/3}$ where $N_T = 30,000$ is the number of terminal nodes in the conducting airway tree.

The algorithm for generating conducting airways can be found in detail elsewhere (Bordas et al., 2015). We will briefly cover it here. Seed points were first clustered based on their nearest terminal branch in the airway tree. The centroid of each cluster was found and a splitting-plane was defined based on the centroid and the node of the parent branch. This splitting-plane was used to divide each seed point cluster into two more clusters. A new branch was generated extending towards each seed point set by 40% of the distance to the centroid. When the seed point set had only one node, or the branch length was less than 2 mm, the branch is classed as a terminal bifurcation and the nearest seed point was deleted. This was repeated until no seed points remained in the tree. Once the tree was grown, a diameter was assigned to each branch segment as described in Supplementary Material Section 3. In our CPF simulation, the image-based airways are used until generation 3 and airway tree generated from the volume-filling algorithm (Bordas et al., 2015) is used to generate a surface from generations 3 to 6, similar to Nousias et al. (2020).

2.5 Deposition simulation configuration

To prepare the airway segmentations of the LUNA16 dataset for deposition simulations, we added a generic mouth-throat geometry obtained from a healthy adult patient (Banko et al., 2015) as done in Williams et al. (2022a) in cases where the upper airways were not imaged. This was done by scaling the surface mesh of the healthy patient’s upper airways such that the diameter at the interface between healthy patient upper airways and the patient’s trachea

reconstructed was the same (Williams et al., 2022a). As jet-like flow structure produced in the upper airways has a key influence on the swirling flow observed in the trachea (Feng et al., 2018), it is important to include at least a generic representation of the upper airways in respiratory CPFD models (Feng et al., 2018; Williams et al., 2022a). For the Southampton/Air Liquide dataset, segmentations of the mouth and throat were obtained by region-growing and joined to the trachea directly without scaling since they were from the same patient.

The deposition simulations in this study were performed using OpenFOAM v6 (Weller et al., 1998) using our custom solver deepLungMPPICFoam (Williams et al., 2022b). Briefly summarised, we solved mass and momentum conservation equations for transport of an incompressible fluid (air). As described in Williams et al. (2022a), our computational fluid dynamics mesh had a spacing of $\Delta = 500 \mu\text{m}$, which gave excellent agreement with flow through 3D printed airways Banko et al. (2015). The fluid turbulence modelling is discussed in detail in Williams et al. (2022c). Particles were tracked in a Lagrangian approach by solving Newton’s equations of motion accounting for drag and gravity, as we showed in our previous study that particle-particle interactions are not influential (Williams et al., 2022a). Particles were deleted when they reached an outlet, and were given a ‘sticking’ condition when they hit the wall boundary. The number of particles tracked was 200,000 with diameter, $d_p = 4 \mu\text{m}$ as used in our previous study (Williams et al., 2022a). In simulations with airways from the LUNA16 dataset, particles were released from the inlet over a period of 0.1 s with initial velocity as 10 m/s to represent a metered-dose inhaler (Liu et al., 2012). More information on simulation configuration for comparing to the Southampton/Air Liquide dataset is given later in this section.

To drive flow in the simulation, we model pressure at the outlet of the 3D domain using a lumped-parameter model based on a resistance-compliance circuit (Oakes et al., 2018). In this model, the pleural pressure required to drive flow during inhalation, p_d , is found by

$$p_d(t) = R_{global}Q(t) + \frac{V(t)}{C_{global}} - p_{atm}, \quad (12)$$

where R_{global} is the global resistance of the circuit, C_{global} is the global capacitance, $Q(t)$ is the flowrate as a function of time, $V(t)$ is the volume of air inhaled as a function of time, and p_{atm} is the atmospheric pressure. As the flow is incompressible, there is no effect of

atmospheric pressure and therefore it is dropped ($p_{atm} = 0$). In simulations of inhalation in airways from the LUNA16 dataset, we model tidal inhalation where the volume of air inhaled to the lung is (Oakes et al., 2018)

$$V(t) = -\frac{1}{2} \left[V_T \cos \left(2\pi \frac{t}{T_B} \right) - V_T \right] \quad (13)$$

where V_T is the tidal volume and T_B is the time for one breath. Oakes et al. (2018) estimated respiratory parameters such as R_{global} , C_{global} , V_T using empirical relationships which are often based on parameters such as age, weight, height, gender. We used the adult parameters from Oakes et al. (2018), $R_{global} = 7 \times 10^{-3}$ cm H₂O s/mL and $C_{global} = 59$ mL/cm H₂O. Pressure at the outlet is then calculated by solving the following equation at each outlet (subscript (i)) at each time-step of our simulation:

$$p_{(i)} = R_{(i)}\phi_{(i)} + \frac{V_{(i)}}{C_{(i)}} + p_d. \quad (14)$$

where $\phi_{(i)}$ is the flowrate (also called ‘flux’) at the outlet computed from the fluid solver. Volume at the outlet is then found by integrating the flux in time. Local compliance and resistance at each outlet are estimated based on

$$C_{(i)} = \frac{A_{(i)}\alpha_{(L)}}{A_{(L)}} C_{global} \quad R_{(i)} = \frac{A_{(L)}}{A_{(i)}\alpha_{(L)}} R_{global} \quad (15)$$

where the $A_{(L)}$ is the sum of areas of all outlets in each lung, L . $\alpha_{(L)}$ is the fraction of (static) volume of a lung to the combined volume of both lungs, which assumes that the volume change during inhalation is a factor of the lung’s volume only (not accounting for localised disease).

By defining Neumann boundary conditions for velocity at the inlets and outlets, the simulation can be unstable. Particularly, when the flow at the outlet is reversed (due to turbulence or local changes in geometry curvature) the flux at the face has a negative sign (creating positive pressure) and can cause numerical divergence (Esmaily Moghadam et al., 2011, 2013). To prevent backflow instabilities, we (i) set maximum pressure at each outlet to zero and minimum flux at each outlet to zero, (ii) use limiters when calculating gradients at outlet faces to prevent small negative fluxes caused by interpolation errors. We also (iii) set

the maximum Courant number to be 0.1 when the flow is developing, and Courant number is set to 0.2 when the flow is developed (quantified by $0.2 < |\sin(2\pi t/T_B)|$).

Simulations of the Southampton/Air Liquide dataset modelled inhalation from a breath-actuated AKITA nebuliser (Conway et al., 2012). This provided a constant inhalation flowrate of 18 L/min. Therefore, we assigned $V(t) = V_T t/T_{inhale}$, where T_{inhale} is the time for one inhalation, $T_{inhale} = T_B/2$. In the experimental study, inhalation duration was varied between 2 and 3.3 s (shallow and deep breathing). However, Fleming et al. (2015) observed no significant difference in regional deposition when changing from shallow to deep inhalation. As our model cannot account for particle motion in the deep lung, we cannot model particles re-entering the central airways from the deep lung during exhalation. Therefore, we used an inhalation duration of 2 s for our simulations to reduce simulation clock-time, since the longer inhalation and breath-hold would only influence particle motion in the deep lung (increasing time for deposition due to diffusion or settling). As the inhalation flowrate is steady and we could not model exhalation, we only modelled one inhalation. This was also done by De Backer et al. (2010) to compare with SPECT-CT data, although they used a steady solver which cannot capture transient flow features.

In the Southampton/Air Liquide experimental study (Conway et al., 2012; Fleming et al., 2015), the particle diameters were $d_p = 3.1$ and $6.05 \mu\text{m}$. To simplify our analysis, we used monodisperse particles as the experimental size distribution was reported as being narrow (Conway et al., 2012). Simulated particles were released from a 5 mm disk in the AKITA mouthpiece, which represented the diameter of the tubing that delivers the particles from the nebuliser. Simulated particles were released with initial velocity of 0 m/s, since the particle timescale was small it would quickly adapt to the local fluid velocity in the mouthpiece ($\tau_p = 28 \mu\text{s}$ for $d_p = 3.1 \mu\text{m}$ and $\tau_p = 108 \mu\text{s}$ for $d_p = 6.05 \mu\text{m}$).

Simulations of flow in airways from the LUNA16 dataset were performed using ROCKS with 28 CPUs each (Intel Xeon Silver), taking between one and three days. Simulations of flow in patient airways from the Southampton/Air Liquide dataset were performed on Oracle cloud with two bare metal compute nodes each with 36 CPUs (BM.optimised3.36, Intel Xeon Gold), taking between one and three days depending on the airway geometry.

2.6 Model evaluation

To evaluate our model, we compared two benchmarks. The first aimed to verify our 2D-3D reconstruction algorithm by comparing reconstructed lung and airway morphological properties, as well as deposition results against ground truth segmentations from the LUNA16 dataset (Section 2.1). The second stage aimed to validate our drug deposition modelling by comparing deposition predictions in airways segmented from CT images with our CNN to *in vivo* deposition data described in Section 2.1 (Conway et al., 2012; Fleming et al., 2015). As our SSAM did not contain training data from the mouth-throat region, we did not test the SSAM against the Southampton/Air Liquide deposition measurements. The Southampton/Air Liquide experimental data contained measurements from deposition in six healthy and six asthmatic patients. Two experiments were performed for each patient, testing varying combinations of aerosol size, breathing pattern (shallow/deep) and inhaled gas (air compared to He/O₂) giving 24 total measurements. We omitted cases with He/O₂ inhaled gas (4/24 experiments). Images in three asthmatic patients had blockages in the upper airways due to tongue positioning, which meant we could not segment the entire upper airways and these patients were omitted (4 of remaining 20 experiments). Additionally, as our model does not account for particle transport in the deep lung or exhalation we omitted the ‘shallow’ inhalation in cases where shallow and deep inhalations were the independent variable compared (4 of remaining 16 cases). Therefore, we performed 12 simulations to evaluate our model against the available experimental data (Conway et al., 2012; Fleming et al., 2015).

To validate the 2D-3D reconstruction algorithm, we reconstructed the lungs against DRRs from the LUNA16 dataset. We evaluated the reconstructed lung space volume against radiologist-verified lung segmentations from the same volumetric CT scan used to create the DRR. During the deep breathing expected when inhaling aerosol drugs, the lung space volume changes significantly. A recent modelling study (Koullapis et al., 2020) applied a lung expansion equal to 70% relative to the initial volume to deform the deep lung’s walls during inhalation. Bearing this in mind we aimed to have a volume error significantly below 70% to produce realistic estimation of lung space volume. We evaluated the quality of our airway SSAM reconstruction by comparing branch diameter to the ground truth, as the diameter will influence flow characteristics such as Stokes number and Reynolds number

in the central airways (Kleinstreuer and Zhang, 2010). Morphological data is compared to ground truth in terms of relative error, and median, upper-quartile and 95th percentiles are assessed.

We evaluated the deposition fraction predictions in SSAM and CNN-based airways compared to simulations using the ground truth airways. Deposition fraction was computed as the number of particles depositing in a region (relative to the total number of inhaled particles). Deposition predictions were assessed based on regional drug deposition fraction (right and left lung), and a local drug deposition concentration metric known as deposition enhancement factor (DEF) (Balashazy et al., 1999; Longest et al., 2006). We calculate this using the number of particles deposited within a fixed distance (1 mm, area $A_{conc} = \pi (1 \text{ mm})^2$) of the central point of each wall face, as used by Dong et al. (2019b). This is made relative to the global deposition by

$$\text{DEF} = \frac{N_{p,deposit}(A_{conc}) / A_{conc}}{N_{p,deposit}(A_{tot}) / A_{tot}}, \quad (16)$$

where $N_{p,deposit}(A)$ is the number of deposited particles in a surface area A . A_{tot} is the total airway surface area. We use this instead of the sum of the deposition efficiencies as the defined areas may overlap, and this therefore prevents particles being counted multiple times towards the global average. To minimise differences in total number of deposited particles or total surface area interfering with quantitative comparison between DEF in the ground truth and SSAM or CNN, we used the ground truth denominator from Equation 16 in DEF calculations in the SSAM and CNN (Dong et al., 2019a). Finally, we validated our framework by quantifying simulated regional deposition (left lung, right lung and mouth-throat airways) against experimental values reported by Conway et al. (2012) and Fleming et al. (2015). Relative errors are expressed in terms of the median, upper quartile and maximum. Correlation between experimental and simulated deposition data is also assessed using the concordance correlation coefficient. All statistical analyses were performed in Python 3.8.




Figure 3: Lung space volume error from SSAM and U-Net reconstructions of patients from LUNA16 dataset. Inset shows U-Net lung space volume error, where the U-Net results were produced using the pretrained U-Net from Hofmanninger et al. (2020).

3 Results

3.1 Morphological assessment on LUNA16 dataset

We found our SSAM to yield a relative lung space volume error of 9.9% median and a 95th percentile of 34.3% on the LUNA16 dataset (Figure 3a). The maximum error for the right lung was 33.64% and 41.66% for the left lung. The concordance correlation coefficient between SSAM lung volume and ground truth lung volume was 0.884 and 0.893, respectively for one and two DRRs used for SSAM fitting (Figure S4). The total volume error (averaged over left and right lungs) had a median value of 9.8% and a 95th percentile of 26.1%. Larger error in the left lung may be due to the presence of the heart interfering with the edge-map, which would cause a poor fit at the inner side of the lung. In contrast, the absolute relative error for the U-Net lung segmentation gave a median value of 1.35% and a 95th percentile of 3% (Figure 3). The U-Net is well-suited to segmenting lungs from CT, due to the high-contrast and well-resolved boundary between dark (air-filled) lung parenchyma and surrounding soft and hard-tissue. The SSAM lung reconstruction yielded a larger maximum error due to the poor contrast between lung and surrounding tissues on some radiographs.

In Figure 4 we compared the output SSAM landmarks against the input X-ray image and lung outline map in two outlier cases and the best case (in terms of lung space volume error, Figure 3). In both cases, we observed that the outer edge of the left lung has incorrectly fitted




Figure 4: Example of fitted SSAM compared to projection and lung outline for two largest lung space volume errors. Yellow markers are output landmarks from the SSAM, and black markers represent the X-ray edges detected by the Canny edge map described in Section 2.2.1. Panels show (a) the largest error, (b) second largest error and (c) the lowest error. In (a) the right lung error was 31% and left lung error was 42.6%. In (b) the right lung error was 15.5% and left lung error was 38.3%. In (c) the right lung error was 1.8% and left lung error was 0.06%.

to a spurious edge introduced to the lung edge map. In the largest error case (Figure 4a), the lower corner of the SSAM landmark point cloud has fitted to the outline of the patient’s torso. In the case with the second largest error (Figure 4b), the left lung outer edge and right lung lower edge have mistakenly fitted to an edge in the soft tissue. In the lowest error case (Figure 4c), there is a spurious edge for the outline of the patient’s torso that slightly interferes with the outer edge of the patient’s right lung. However, this is located very near the lung itself as the patient has a narrow torso, which therefore appears to cause a minor influence on the SSAM fitting.

Diameter predictions did appear to be sensitive to outliers, as the maximum diameter error was above 30% in the main bronchi (Figure 5). We compare airway diameter error with a single DRR (anterior-posterior plane projection) and two DRRs (anterior-posterior and sagittal plane projections). The maximum trachea diameter error was 20.4% for one projection and 21.9% for two projections, which shows the trachea is less sensitive to outliers than the main bronchi due to higher visibility on an X-ray image. The absolute trachea diameter error had a median value of 8.5% and 7.8% for one and two projections, respectively. The main bronchi maximum error was 26.9% and 33.2% for one and two projections, and the upper quartile error was 13.6% and 12.4%. The concordance correlation coefficient was slightly higher for two projections than one (0.71 compared to 0.657, Figure S5). There

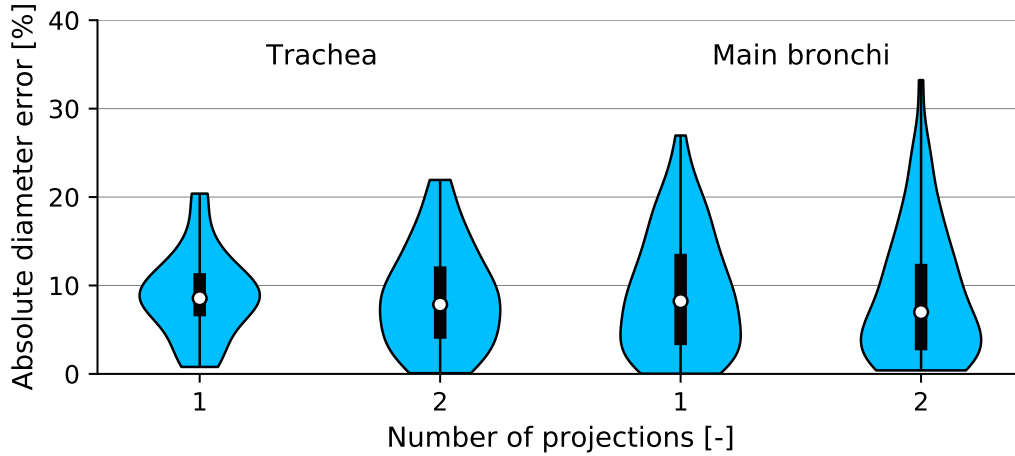


Figure 5: Comparison of airway diameter errors predicted by our SSAM compared to the ground truth segmentations. We show the influence of including an additional projection on diameter error in the trachea and main bronchi.

was no statistically significant difference between the absolute error when comparing SSAM results with one and two projections (Wilcoxon signed-rank test $p = 0.6$ and $p = 0.69$ for the trachea and main bronchi, respectively).

To evaluate the effect of imaging error propagation into distal airway diameters, we evaluated the diameter generated from the distal airway generation algorithm from the ground truth compared to the SSAM and CNN (Figure 6). Here, the ground truth refers to results from the airway generation algorithm using the central airway tree from the ground truth segmentations. Due to the dependence on the diameter in the central airways, the CNN also shows mixed performance. The median absolute error in mean airway diameter per generation was 24.2% (similar to the SSAM with parent model, Figure S6). Additionally, the maximum absolute error was 77.6%, compared to 76% for the SSAM. The CNN airways show improved predictions of the diameter standard deviation (Figure 6), as the maximum error in standard deviation was 125.5% for the SSAM and 69.6% for the CNN.

3.2 Deposition analysis

In this section, we first compare deposition predictions from CFPD simulations in the LUNA16 dataset. We compare predicted deposition in ground truth airways with deposition predictions in airways obtained from SSAM and CNN segmentations. Following this, we then compare our simulated deposition from airways segmented by the CNN with experimental




Figure 6: Generated airway diameter for our image-processing modalities on two cases from the LUNA16 dataset. Markers show the normalised mean diameter per generation. Error bars show $\pm 1S.D.$ where $S.D.$ is the standard deviation.

results from the Southampton/Air Liquide dataset.

Errors in deposition fraction from simulations of flow in airways from the LUNA16 dataset are shown in Figure 7. Comparing simulations in the CNN with the ground truth, the median absolute error was 6.5%, with a 95% confidence interval of $-3.3 \pm 12.1\%$ (Figure 7a). The Pearson correlation coefficient between ground truth and CNN deposition results was 0.88. Deposition in the SSAM airway had a similar median absolute error as the CNN deposition (8%, Figure 7b). SSAM deposition error was more widely distributed, as the 95% confidence interval was $-3.5 \pm 20.5\%$. The Pearson correlation coefficient between deposition in the SSAM with the ground truth was 0.6. The error range for deposition in the SSAM was similar to diameter error (Figure 5) which had a confidence interval of $-2.1 \pm 21.7\%$.

To evaluate the effect of SSAM reconstruction on local deposition hotspots, we computed the deposition enhancement factor (DEF). The deposition hotspots are highly sensitive to local flow and geometry (Farghadan et al., 2020; Williams et al., 2022a), which itself is susceptible to geometric differences by image acquisition and image processing (MacDonald et al., 2020). We can see that even for the case with the lowest deposition error (Figure 8), the deposition hotspots in the SSAM do not match the ground truth. Particularly, in the ground truth and CNN segmented airways (Figure 8a,c), the drug is noticeably more dispersed than in the SSAM (Figure 8b). Less dispersed hotspots in the SSAM is likely due to additional turbulence generated in the CT-based airways which is not present in the SSAM, since it con-

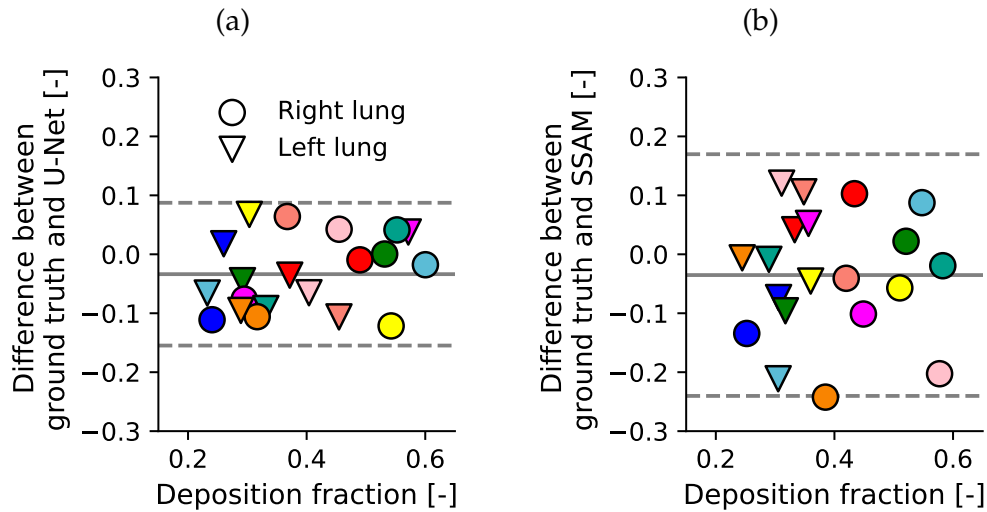


Figure 7: Bland-Altman plot comparing drug deposition in each lung predicted in ground truth airways, compared to airways reconstructed by our SSAM. Number of patients in sample was 9. Each marker represents one lung, and each colour represents a different patient.

tains a reduced number of principal components. A lower number of principal components essentially simplifies the geometry and removes intricate details that may onset turbulence such as cartilagenous rings in the trachea (Zhang and Finlay, 2005; Russo et al., 2008).

Deposition hotspots in the CNN-based airways appear qualitatively similar to the ground truth (Figure 8a,c). The surface area with $DEF > 1$ at the first bifurcation in the ground truth is 17% larger than the CNN. However, the intensity of the hotspots is significantly different since the average DEF at the first bifurcation was 2.49 for the ground truth (maximum value 400), compared to 0.93 for the CNN (maximum value 225). At the first bifurcation, the mean DEF in our SSAM was 1.06 (maximum value 225). The surface area with $DEF > 1$ at the first bifurcation was 2.16x larger in the ground truth than the SSAM. Interestingly, the mean DEF at the first bifurcation in the SSAM and U-Net are similar (1.06 compared to 0.93 for SSAM and U-Net, respectively). However, there is a significant difference in the surface area covered by the deposition hotspots here (Figure 8).

Finally, we performed simulations of nebulised aerosol inhalation in healthy and asthmatic patient cases of the Southampton/Air Liquide dataset (Conway et al., 2012; Fleming et al., 2015). The maximum deposition absolute error averaged over both lungs was 12.1%, which belonged to the case with most obstructed upper airway images shown in Figure S1a (case H02 with small particles). The median, 75th percentile and maximum absolute

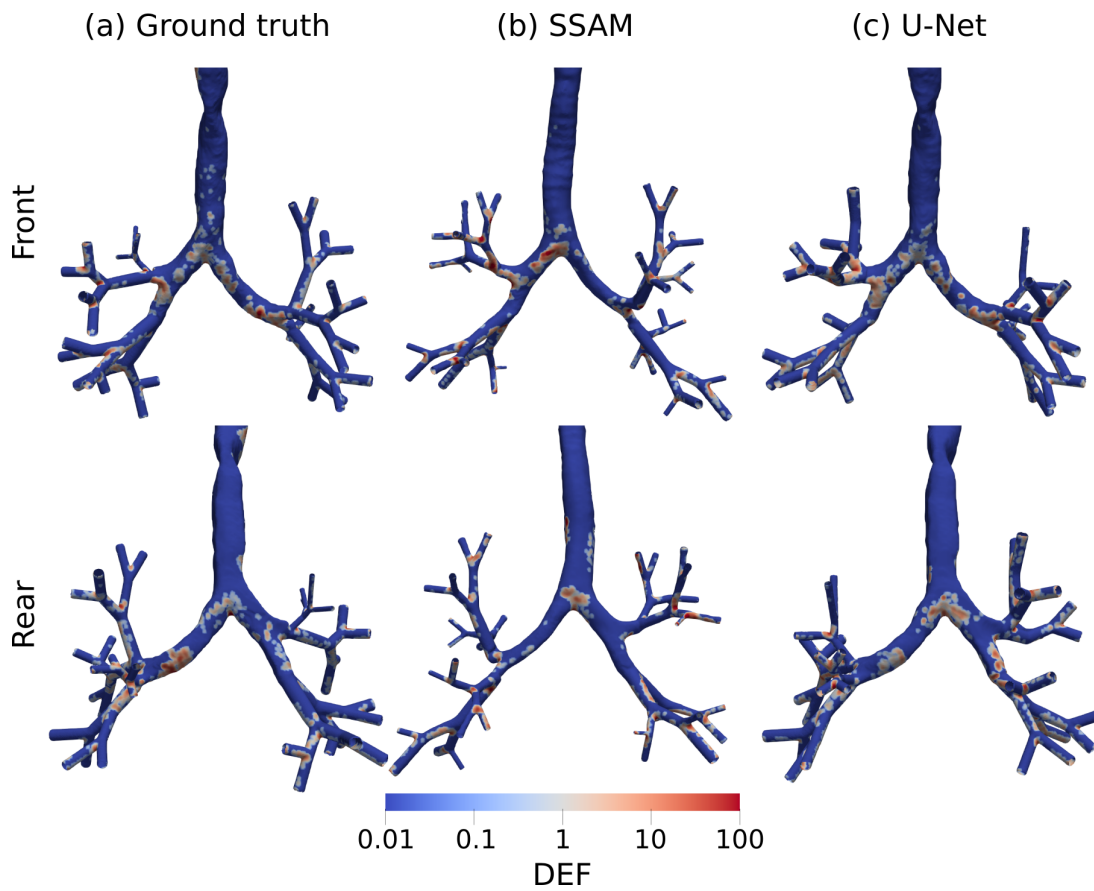


Figure 8: Deposition enhancement factor for simulations in airways constructed using the developed SSAM and U-Net, compared to simulations in a ground truth segmentation. Columns show (left) ground truth, (middle) SSAM, and (right) U-Net segmented airways. Rows show (top) front and (bottom) rear view of the airways.

error was 4.9%, 7.9% 13.1%, respectively. The concordance correlation coefficient between simulated and experimental lung deposition measurements was 0.43. The mouth-throat deposition had a concordance correlation coefficient of 0.525 (Figure 9b). The maximum mouth-throat deposition error was also in case H02 with small particles (27.2%). The linear best-fit line shows that the cases with high error in the mouth-throat region causes the distribution to depart from the 45 ° line (Figure 9a,b). We observed that mouth-throat deposition was best in the two cases where the glottis did not require manual cleanup (grey and yellow markers in Figure 9b). The error in these cases was 1.3% and 3.1%. The median and 75th percentile mouth-throat deposition errors were 5.7% and 11%, respectively. These results show that our models agree well with *in vivo* deposition data, particularly in the absence of imaging artefacts (Figure 9).

In an attempt to understand how our simulations would perform in a perfect case with no imaging difficulties in the upper airways, we partially corrected the simulation and experimental data by normalising the deposition fraction by $1 - DF_{mouth}$, where DF_{mouth} is the deposition fraction in the mouth (Figure 9c). This makes deposition relative to number of particles entering the trachea, rather than the total number of particles entering the mouth. We apply this only to cases where the experimental lung deposition fraction was 5% larger than the simulated lung deposition (only outlier cases above the 45° line), as these were cases where the mouth-throat deposition was severely over-estimated (for example, green markers in Figure 9). Once these problematic cases were corrected, we see an improvement in the correlation between experimental and simulated data, as the concordance correlation coefficient increases from 0.432 to 0.810 (Figure 9a,c). This can be seen quantitatively by the fitted linear regression line nearly matching the 45° line. The resultant r^2 value is 3.3 times larger, at a value of $r^2 = 0.699$. The median, 75th percentile and maximum errors all decreased to 2.12%, 4.14% and 8.6%, respectively. This idealised test shows that the physical model is working well, but manually cleaning the segmentations around the glottis can introduce uncertainty that lowers correlation with experimental data (Figure 9a,c). Additionally, this idealised test points to an avenue for future development of image-processing tools for this region.

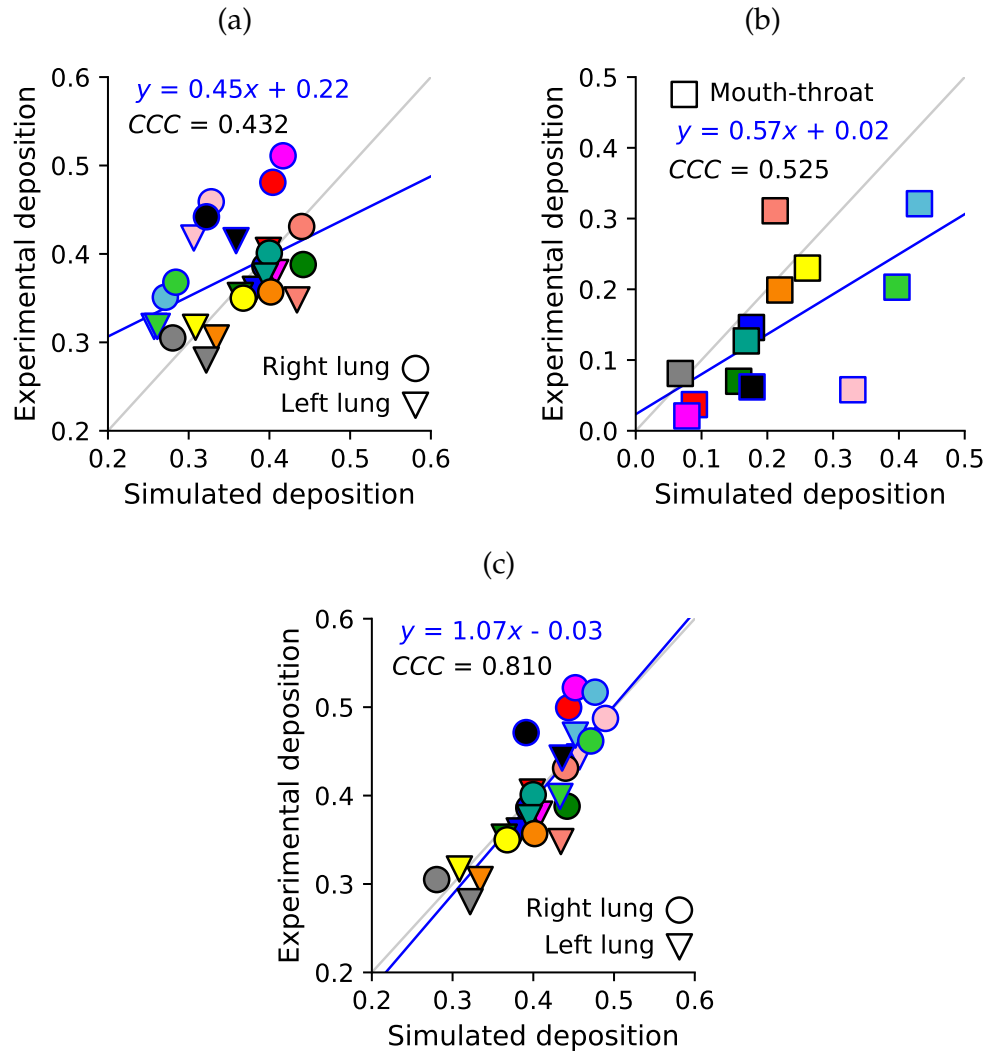


Figure 9: Comparison of simulated deposition against experimental measurements from Conway et al. (2012). Panels show (a) right and left lung deposition, (b) mouth-throat deposition, and (c) the partially corrected right and left lung deposition. Cases which are adjusted by correcting over-predicted mouth-throat deposition are shown with a blue border on the marker. CCC is the concordance correlation coefficient. The linear regression lines shown have (a) $r^2 = 0.213$, (b) $r^2 = 0.36$, (c) $r^2 = 0.699$. The cases requiring no manual cleanup of the glottis have grey and yellow markers.

4 Discussion

In this study, we aimed to develop an image-processing framework to automatically generate personalised computational respiratory systems from volumetric CT and planar X-ray image(s). To that end, we developed a statistical shape and appearance model to parameterise the variation in 3D lung and airway shapes across a population. An algorithm was proposed to reconstruct 3D lung and airway geometries, by iteratively adjusting shape parameters based on a given X-ray image. We also trained a convolutional neural network (CNN) based on the U-Net and ENet architectures (Ronneberger et al., 2015; Paszke et al., 2016), which allow us to automatically segment airways from volumetric CT scans. We evaluated the ability of our proposed framework to reproduce morphological properties and deposition results that were obtained from the ground truth airway/lung geometries. The CNN-based airways showed good agreement with the ground truth on all metrics. Our SSAM showed good agreement with the ground truth, although it was sensitive to outliers in predictions of morphological properties due to the low amount of information available in a planar X-ray image. Finally, we aimed to validate our proposed framework against experimental *in vivo* deposition measurements (Conway et al., 2012; Fleming et al., 2015). After cleaning the upper airway segmentations to account for glottis motion artefacts, we obtained good agreement with the experimental data.

Despite imaging uncertainties in the upper airway geometry, we obtained good agreement with experimental deposition measurements (median deposition error 4.9%, Figure 9). Available *in vivo* experimental studies of drug deposition are limited due to the difficult and expensive nature of the experiments. Therefore, most modelling studies with validation compare to *in vitro* experiments, such as flow or deposition measurements in airway casts (Williams et al., 2022a; Koullapis et al., 2018b; Holbrook and Longest, 2013; Longest et al., 2012). Alternatively, some CPFD validation studies have used varying forms of experimental *in vivo* data (Tian et al., 2015; De Backer et al., 2010; Oakes et al., 2013, 2015). Tian et al. (2015) compared their geometrically-simplified model to 2D scintigraphy data at varying particle sizes. De Backer et al. (2010) compared patient-specific models with SPECT-CT data, but only for tracer particles ($d_p \approx 1 \mu\text{m}$). Additionally, their outlet boundary conditions uses a fixed flowrate based on the regional lung space volume change from CT images at full in-

halation and exhalation. In contrast, the outlet pressure in our study is approximated with a model based on lung space volume fraction, resistance and compliance (Section 2.5). Our approach has the benefit that it requires less radiation (only one CT or X-ray image required) to model regional ventilation to each lung. Other *in vivo* validation studies have compared to 3D deposition measurements in rats (Oakes et al., 2013, 2015), although the Reynolds number and Stokes number will change in these cases since the airway diameter is smaller than humans. Therefore, by validating our full patient-specific modelling and imaging approach against 3D *in vivo* data at varying particle sizes, our work represents a key development in advancing towards incorporating deposition models into clinical care.

Morphological and deposition results produced by the developed SSAM had a wide distribution, as the SSAM was sensitive to outliers. We observed that outliers occurred in cases where the outline of the lung on the X-ray image was noisy or low quality (Figure 4). However, the SSAM largely provided good agreement with the ground truth, as median absolute error was under 10% for the lung space volume (Figure 3), airway diameter (Figure 5), and lung deposition (Figure 7b). We also observed that the maximum error for airway diameter (32%) was less than the maximum error for regional deposition (24%), which shows that morphological errors in the SSAM do not propagate into significantly worse regional deposition results despite the nonlinear influence of diameter error on local velocity, Reynolds number and Stokes number on deposition.

When assessing local deposition hotspots, we observed significantly different deposition patterns in the SSAM compared to the ground truth (Figure 8a,b). Therefore, for exact predictions of deposition hotspots, the SSAM is not a suitable approach and a CNN segmentation from CT should be used. This occurs since hotspots are highly influenced by near-wall flow structures. This shows the complex geometrical features of the airways that create secondary flows responsible for deposition cannot be reconstructed from the limited information available in a planar X-ray image. Our SSAM shows similar performance to the SSAM developed by Väänänen et al. (2015) to reconstruct femora from 2D radiographs, as the average errors were below 10%, but the SSAM had difficulty predicting local morphological properties. The quality of our reconstruction from 2D images may be improved by using a neural network (transformer) architecture to find the optimal shape parameters (**b**)

from the X-ray image, as this approach was shown by Shen et al. (2019) to reconstruct a full volumetric CT from a series of X-ray images.

By analysing the lung space volume error outliers, we observed that the input lung edge map appeared to create some difficulties due to spuriously introduced edges (Figure 4). Li et al. (2016) proposed an approach to extract lung-field outlines from chest X-rays using SSAMs. Using a SSAM instead of a general edge detection algorithm constrains the lung outline with the SSAM modes, and ensures no non-lung edges are highlighted to potentially interfere with the fitting. Alternatively, fitting of the SSAM could be improved by replacing the global loss function parameters (model coefficients, C in Equation 2) with patient-specific values. This could be obtained by determining the optimal parameters for each patient in the training database, through the same Gaussian process regression described in Section 2.2.2. The input chest X-ray may then be passed through a convolutional neural network with four output neurons (one for each C in Equation 2). These minor alterations may improve fitting the loss function given in Equation 2 and yield a more robust SSAM algorithm.

Error between experimental and simulated deposition in CNN-segmented airways in the LUNA16 dataset gave a 95% confidence interval of $-3.3 \pm 12.1\%$ (Figure 7a). This is likely due to errors in distal airway diameter (Figure 6) that influenced deposition in the distal airways and particles reaching the deep lung. The distal airway diameter may also have influenced CNN-segmented airway deposition through changing the outlet diameter and the computed outlet boundary condition through Equation 15. The ability of the CNN to segment distal airways was limited by the low resolution of LUNA16 datasets (mean slice thickness 1.6 mm), which is much lower than the high resolution CT used in the Southampton/Air Liquide dataset (mean slice thickness 0.5 mm) (Conway et al., 2012; Fleming et al., 2015). We are also limited by the quality of ground truth obtained by our semi-automatic algorithm which was verified by a radiologist to have no leakage, but suffered from a limited number of branches segmented. This was due to the low image resolution, which can create significant difficulty in obtaining a full airway tree without leakage (Tschirren et al., 2005a; Lo et al., 2012). The CNN training could be improved to mitigate this by masking voxels belonging to trachea and main bronchi, which forces the model to focus on learning to segment small airways (Garcia-Uceda et al., 2021).

To move towards incorporating this framework in a clinical decision-making workflow, the bottleneck of simulation time is the primary obstacle to overcome. Clock-time for simulations in this study ranged between 1-7 days, which is not suitable for clinical use. To deliver real-time prediction of clinical quantities of interest, geometric parameters (such as SSM weights) and flow variables can be used to fit machine learning models to simulation data to provide predictions of deposition within minutes (Hoeijmakers et al., 2020, 2021; Su et al., 2020). Morales Ferez et al. (2021) showed that wall-shear stress could be predicted from patient-specific vascular meshes with minimal pre-processing using geometric deep learning, which has been developed for analysing non-Euclidean geometries such as graphs or meshes (Bronstein et al., 2017). This approach could be applied to the respiratory system to evaluate properties such as deposition hotspots (Figure 8) based on a mesh obtained with our segmentation algorithm, without the additional step of fitting to a SSM as done by Hoeijmakers et al. (2020) for the aortic valve. An alternative to using classic purely data-driven machine learning models to predict deposition is to constrain the neural network fitting with some physical constraints such as conservation of mass and momentum (Raissi et al., 2018, 2019; Alber et al., 2019; Karniadakis et al., 2021). This approach generates full velocity and pressure fields based on the physical constraints and training data from similar flow configurations. However, it is unclear if the physics-informed neural network approach would work in the geometrically complex airway tree, since its application has been limited to simple flow configurations such as bluff body flow, lid-driven cavity or pipe flow (Raissi et al., 2018; Yin et al., 2021; Jagtap et al., 2020). Our developed imaging and simulation framework could be used to generate training data for machine learning approaches, as well as to benefit from the speed improvements created by machine learning for CFD. Combining these modelling approaches is crucial to feasibly incorporate patient-specific models in clinical practice.

4.1 Limitations

A limitation of this study was the difficulty in capturing the upper airway anatomy (Supplementary material Section 1). The glottis can dilate and narrow during inhalation and exhalation, respectively (Brancatisano et al., 1983; Scheinherr et al., 2015), which created dif-

difficulties in capturing CT images of the glottis (Conway et al., 2012). Changes in the shape and area of the airways here can cause increased deposition in the glottis (Zhao et al., 2020), but also downstream due to the flow separation at the expansion which creates an unsteady wake whose structure will likely vary based on patient anatomy (Luo and Pedley, 1996; Cisonni et al., 2010; Bhardwaj et al., 2022). Therefore, our attempt at cleaning up the glottis cannot perfectly represent instantaneous patient-specific flow patterns observed *in vivo*. Despite this, we observed good agreement with experimental lung and mouth-throat deposition measurements (Figure 9). Future efforts should be focussed towards combining our developments with dynamic upper airway models (Zhao et al., 2020; Bhardwaj et al., 2022), to enable modelling the complex biomechanical fluid-structure interaction of respiration. However, to perform fluid-structure interaction simulations of these cases still requires a reliable baseline airway model, which requires manual cleanup of outlier cases shown in Figure 9. An improved approach to manually cleaning the glottis segmentations may be to determine the missing airway structure based on statistical shape models, as done for segmenting obstructed bronchi by Irving et al. (2011). Therefore, a template mesh could be morphed to approximate the glottis shape based on the shape of the nearby upstream and downstream airways (the throat and trachea). Combining this improved segmentation cleanup approach with dynamic mesh simulations in the upper airways is likely to produce a robust approach for modelling upper airway dynamics.

An additional limitation is the use of a healthy patient mouth-throat geometry (Banko et al., 2015) as a representative generic geometry for the upper airways in the LUNA16 dataset. Due to unavailable imaging data above the trachea in most clinical CT datasets, it was not possible to segment this from CT with our CNN, or include it as part of our SSAM to reconstruct from X-ray images. This of course limits the understanding of how well our imaging approach can capture the upper airways and flow patterns specific to a patient. The effect of this would be negligible on the CNN segmentation, as the upper airways are easily segmented with a region-growing approach (Garcia-Uceda et al., 2021). It is unclear how well the SSAM could reconstruct the mouth-throat from a patient X-ray image, as there is no existing SSM that shows the main modes of variance in this part of the airways. However, literature agrees that the glottis cross-sectional area is the most significant

upper airway morphological parameter influencing deposition (Xi et al., 2016; Feng et al., 2018). Moreover, realistic mouth-throat geometries were shown to give similar regional deposition predictions as a simplified mouth-throat geometry with an elliptical cross-section when $d_p \leq 10 \mu\text{m}$ (Xi et al., 2016). This suggests that it would not be essential for a SSAM to capture the exact shape of the mouth-throat, but an accurate prediction of the glottal cross sectional area is required. With a good estimation of glottal area, the SSAM could produce deposition that agrees with the deposition as shown for our generic healthy patient mouth-throat (Figure 7).

Our deep lung model was a 0D approximation of deep lung mechanics (Oakes et al., 2015, 2018). Using this 0D model as an outlet boundary condition for the 3D domain (3D-0D model) has been shown to accurately predict lobar deposition in healthy and emphysematous rats (Oakes et al., 2015). In this 0D model, particles reaching the outlets are deleted and classified as depositing in the deep lung. Oakes et al. (2017) developed a 1D model for fluid and particle transport in distal airways (3D-1D model), which tracks particle concentration in the distal airways. This 1D formulation showed minor improvements in agreement with experimental lobar deposition measurements over the 0D formulation (Oakes et al., 2017). A key benefit of the 1D formulation is that particles can re-enter the 3D domain during exhalation as it is known how many particles are still floating in the distal lung (Oakes et al., 2017). An inability to model particles re-entering the domain during exhalation is a limitation of our model, as we cannot compare the fraction of drug exhaled with experimental measurements (Conway et al., 2012; Fleming et al., 2015). In the majority of experiments, the exhaled fraction was below 10%, meaning this has only a minor influence on the validity of our model for inhalation studies. Incorporating a 1D model for particle and flow transport in distal airways would enable further comparison to SPECT-CT data, as we could assess if our model captures how deep particles penetrate into the lung, rather than only regional deposition given here (Figure 9).

Our 0D deep lung model used global resistance and compliance metrics which were based on empirical correlations from experimental measurements (Oakes et al., 2018; D'angelo et al., 1989). The resistance and compliance per outlet was then based on the outlet area and the lung's volume fraction (relative to combined volume). Oakes et al. (2018) used this pro-

cedure as a initial condition, and then determined the true resistance (accounting for flow resistance in the 3D domain) by iteratively computing the pressure drop between the trachea and each outlet and adjusting the resistance until convergence. This would require many simulations for each patient to determine the exact distribution of flow and particles delivered to each individual outlet in the domain. As we are only interested in regional deposition (right and left lung) at this stage to validate our imaging and model predictions, this was not necessary. Therefore, we have not accounted for 3D domain flow resistance as it has no effect on the distribution of particles between right and left lung, which is dominated by the lobe volume fraction, $\alpha_{(L)}$, and airway morphology. In future studies where exact local deposition is required, the iterative simulation approach of Oakes et al. (2018) may be improved by combining full 3D-0D simulations with a low-fidelity surrogate model to decrease the number of time-consuming simulations (Perdikaris and Karniadakis, 2016).

5 Conclusion

We have developed a rapid image-processing and mathematical modelling pipeline to enable patient-specific predictions of drug deposition with data as sparse as a single chest X-ray image. This approach segmented lungs and airways from volumetric CT data with high accuracy using convolutional neural networks (Ronneberger et al., 2015; Paszke et al., 2016; Hofmanninger et al., 2020). We also developed an approach to extract lung and airway geometries from chest X-rays using a statistical shape and appearance model that iteratively adjusted shape parameters based on the outline and gray-value distribution of an unseen X-ray image. This approach was shown to reconstruct patient respiratory morphologies in good agreement with the ground truth data, with the exception of a few outliers. Deposition in airways from both automated approaches agreed well with deposition in ground truth airways obtained by semi-automatic segmentation. The SSAM reconstruction failed to reproduce local deposition patterns present in the ground truth, which we expect is due to morphological differences and the reduced number of principal components creating a smoother airway. Finally, the imaging and modelling framework was compared to experimental *in vivo* measurements (Conway et al., 2012; Fleming et al., 2015). The quality of images of the upper airway (mouth-throat) was mixed and required cleanup (Conway et al.,

2012), which introduced uncertainty. Despite this uncertainty in the mouth-throat geometry, we achieved good agreement with experimental regional deposition. To enable future integration of physical models into healthcare settings and e-health frameworks, the ability to generate patient-specific respiratory systems and deposition predictions automatically from sparse data is crucial. Predicted deposition information could be used as part of the clinical development of new inhaled therapeutics, and reduce the need for expensive and irradiating *in vivo* deposition imaging. Our developed and experimentally validated framework is an essential step towards clinical implementation of patient-specific modelling, which will allow for automated and reliable predictions of patient therapeutic response to an aerosol drug.

Acknowledgements

Simulations reported in this study were performed on Oracle cloud computing platform, funded by Open Clouds Research Environments (OCRE) 'Cloud Funding for Research'. JW was funded by a 2019 PhD Scholarship from the Carnegie-Trust for the Universities of Scotland. The *in vivo* deposition data used in this study was obtained from a project sponsored by Air Liquide. The authors thank Daniel Bustamante for his work in performing LUNA16 airway segmentations.

References

- M. Alber, A. Buganza Tepole, W. R. Cannon, S. De, S. Dura-Bernal, K. Garikipati, G. Karniadakis, W. W. Lytton, P. Perdikaris, L. Petzold, et al. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ digital medicine*, 2(1):1–11, 2019.
- T. Albrecht, M. Lüthi, T. Gerig, and T. Vetter. Posterior shape models. *Medical image analysis*, 17(8):959–973, 2013.
- P. Angelikopoulos, C. Papadimitriou, and P. Koumoutsakos. Bayesian uncertainty quantifi-

- cation and propagation in molecular dynamics simulations: a high performance computing framework. *The Journal of chemical physics*, 137(14):144103, 2012.
- N. Baka, B. L. Kaptein, M. de Bruijne, T. van Walsum, J. Giphart, W. J. Niessen, and B. P. Lelieveldt. 2D–3D shape reconstruction of the distal femur from stereo X-ray imaging using statistical shape models. *Medical image analysis*, 15(6):840–850, 2011.
- I. Balashazy, W. Hofmann, and T. Heistracher. Computation of local enhancement factors for the quantification of particle deposition patterns in airway bifurcations. *Journal of Aerosol Science*, 30(2):185–203, 1999.
- A. Banko, F. Coletti, D. Schiavazzi, C. Elkins, and J. Eaton. Three-dimensional inspiratory flow in the upper and central human airways. *Experiments in Fluids*, 56(6):117, 2015.
- S. Bhardwaj, P. Koullapis, S. C. Kassinos, and J. Sznitman. Fate of inhaled aerosols under the influence of glottal motion in a realistic in silico human tracheobronchial tree model. *European Journal of Pharmaceutical Sciences*, 173:106172, 2022.
- R. Bordas, C. Lefevre, B. Veeckmans, J. Pitt-Francis, C. Fetita, C. E. Brightling, D. Kay, S. Siddiqui, and K. S. Burrowes. Development and analysis of patient-based complete conducting airways models. *PloS one*, 10(12):e0144105, 2015.
- T. Brancatisano, P. Collett, and L. Engel. Respiratory movements of the vocal cords. *Journal of Applied Physiology*, 54(5):1269–1276, 1983.
- M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- J. Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6):679–698, 1986.
- J. C. Carr, W. R. Fright, and R. K. Beatson. Surface interpolation with radial basis functions for medical imaging. *IEEE transactions on medical imaging*, 16(1):96–107, 1997.
- I. Castro-Mateos, J. M. Pozo, T. F. Cootes, J. M. Wilkinson, R. Eastell, and A. F. Frangi. Statistical shape and appearance models in osteoporosis. *Current osteoporosis reports*, 12(2):163–173, 2014.

- J. J. Cerrolaza, M. L. Picazo, L. Humbert, Y. Sato, D. Rueckert, M. Á. G. Ballester, and M. G. Linguraru. Computational anatomy for multi-organ analysis in medical imaging: A review. *Medical image analysis*, 56:44–67, 2019.
- Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- J. Cisonni, A. Van Hirtum, X. Y. Luo, and X. Pelorson. Experimental validation of quasi-one-dimensional and two-dimensional steady glottal flow models. *Medical & biological engineering & computing*, 48(9):903–910, 2010.
- K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, et al. The cancer imaging archive (tcia): maintaining and operating a public information repository. *Journal of digital imaging*, 26(6):1045–1057, 2013.
- A. Comelli, N. Dahiya, A. Stefano, V. Benfante, G. Gentile, V. Agnese, G. M. Raffa, M. Pilato, A. Yezzi, G. Petrucci, et al. Deep learning approach for the segmentation of aneurysmal ascending aorta. *Biomedical Engineering Letters*, 11(1):15–24, 2021.
- A. Comerford, C. Förster, and W. A. Wall. Structured tree impedance outflow boundary conditions for 3D lung simulations. *Journal of biomechanical engineering*, 132(8):081002, 2010.
- J. Conway, J. Fleming, C. Majoral, I. Katz, D. Perchet, C. Peebles, L. Tossici-Bolt, L. Collier, G. Caillibotte, M. Pichelin, et al. Controlled, parametric, individualized, 2-D and 3-D imaging measurements of aerosol deposition in the respiratory tract of healthy human subjects for model validation. *Journal of aerosol science*, 52:1–17, 2012.
- T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995.
- T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Transactions on pattern analysis and machine intelligence*, 23(6):681–685, 2001.
- E. D’angelo, E. Calderini, G. Torri, F. Robatto, D. Bono, and J. Milic-Emili. Respiratory me-

- chanics in anesthetized paralyzed humans: effects of flow, volume, and time. *Journal of Applied Physiology*, 67(6):2556–2564, 1989.
- J. W. De Backer, W. G. Vos, S. C. Vinchurkar, R. Claes, A. Drollmann, D. Wulfrank, P. M. Parizel, P. Germonpré, and W. De Backer. Validation of computational fluid dynamics in CT-based airway models with SPECT/CT. *Radiology*, 257(3):854–862, 2010.
- J. Dong, Y. Shang, L. Tian, K. Inthavong, D. Qiu, and J. Tu. Ultrafine particle deposition in a realistic human airway at multiple inhalation scenarios. *International journal for numerical methods in biomedical engineering*, page e3215, 2019a.
- J. Dong, L. Tian, and G. Ahmadi. Numerical assessment of respiratory airway exposure risks to diesel exhaust particles. *Experimental and Computational Multiphase Flow*, 1(1):51–59, 2019b.
- J. Dong, Y. Yang, and Y. Zhu. Recent advances in the understanding of alveolar flow. *Biomecrofluidics*, 16(2):021502, 2022.
- J. Dworzak, H. Lamecker, J. von Berg, T. Klinder, C. Lorenz, D. Kainmüller, H. Seim, H.-C. Hege, and S. Zachow. 3D reconstruction of the human rib cage from 2D projection images using a statistical shape model. *International journal of computer assisted radiology and surgery*, 5(2):111–124, 2010.
- M. Esmaily Moghadam, Y. Bazilevs, T.-Y. Hsia, I. E. Vignon-Clementel, and A. L. Marsden. A comparison of outlet boundary treatments for prevention of backflow divergence with relevance to blood flow simulations. *Computational Mechanics*, 48(3):277–291, 2011.
- M. Esmaily Moghadam, I. E. Vignon-Clementel, R. Figliola, A. L. Marsden, M. of Congenital Hearts Alliance (MOCHA) Investigators, et al. A modular numerical method for implicit 0d/3d coupling in cardiovascular finite element simulations. *Journal of Computational Physics*, 244:63–79, 2013.
- A. Farghadan, K. Poorbahrami, S. Jalal, J. M. Oakes, F. Coletti, and A. Arzani. Particle transport and deposition correlation with near-wall flow characteristic under inspiratory air-flow in lung airways. *Computers in biology and medicine*, 120:103703, 2020.

- Y. Feng, J. Zhao, C. Kleinstreuer, Q. Wang, J. Wang, D. H. Wu, and J. Lin. An in silico inter-subject variability study of extra-thoracic morphology effects on inhaled particle transport and deposition. *Journal of Aerosol Science*, 123:185–207, 2018.
- Y. Feng, J. Zhao, H. Hayati, T. Sperry, and H. Yi. Tutorial: Understanding the transport, deposition, and translocation of particles in human respiratory systems using computational fluid-particle dynamics and physiologically based toxicokinetic models. *Journal of Aerosol Science*, 151:105672, 2021.
- L. Ferrarini, H. Olofsen, W. M. Palm, M. A. Van Buchem, J. H. Reiber, and F. Admiraal-Behloul. Games: growing and adaptive meshes for fully automatic shape modeling and analysis. *Medical image analysis*, 11(3):302–314, 2007.
- J. Fleming, J. Conway, C. Majoral, L. Tossici-Bolt, I. Katz, G. Caillibotte, D. Perchet, M. Pichelin, B. Muellinger, T. Martonen, et al. The use of combined single photon emission computed tomography and x-ray computed tomography to assess the fate of inhaled aerosol. *Journal of aerosol medicine and pulmonary drug delivery*, 24(1):49–60, 2011.
- J. Fleming, J. Conway, C. Majoral, I. Katz, G. Caillibotte, M. Pichelin, S. Montesantos, and M. Bennett. Controlled, parametric, individualized, 2-D and 3-D imaging measurements of aerosol deposition in the respiratory tract of asthmatic human subjects for model validation. *Journal of aerosol medicine and pulmonary drug delivery*, 28(6):432–451, 2015.
- A. Garcia-Uceda, R. Selvan, Z. Saghir, H. Tiddens, and M. de Bruijne. Automatic airway segmentation from computed tomography using robust and efficient 3-D convolutional neural networks. *Scientific Reports*, 2021.
- O. Gozes and H. Greenspan. Lung Structures Enhancement in Chest Radiographs via CT based FCNN Training. In *Image Analysis for Moving Organ, Breast, and Thoracic Images*, pages 147–158. Springer, 2018.
- L. Grassi, N. Hraiech, E. Schileo, M. Ansaloni, M. Rochette, and M. Viceconti. Evaluation of the generality and accuracy of a new mesh morphing procedure for the human femur. *Medical engineering & physics*, 33(1):112–120, 2011.

- B. Haefeli-Bleuer and E. R. Weibel. Morphometry of the human pulmonary acinus. *The Anatomical Record*, 220(4):401–414, 1988.
- A. A. Hagberg, D. A. Schult, and P. J. Swart. Exploring network structure, dynamics, and function using networkx. In G. Varoquaux, T. Vaught, and J. Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.
- T. Heimann and H.-P. Meinzer. Statistical shape models for 3D medical image segmentation: a review. *Medical image analysis*, 13(4):543–563, 2009.
- M. Hoeijmakers, I. Waechter-Stehle, J. Weese, and F. Van de Vosse. Combining statistical shape modeling, CFD, and meta-modeling to approximate the patient-specific pressure-drop across the aortic valve in real-time. *International journal for numerical methods in biomedical engineering*, 36(10):e3387, 2020.
- M. Hoeijmakers, W. Huberts, M. Rutten, and F. van de Vosse. The impact of shape uncertainty on aortic-valve pressure-drop computations. *International Journal for Numerical Methods in Biomedical Engineering*, 37(10):e3518, 2021.
- P. Hofemeier and J. Sznitman. The role of anisotropic expansion for pulmonary acinar aerosol deposition. *Journal of biomechanics*, 49(14):3543–3548, 2016.
- J. Hofmanninger, F. Prayer, J. Pan, S. Röhrich, H. Prosch, and G. Langs. Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem. *European Radiology Experimental*, 4(1):1–13, 2020.
- L. T. Holbrook and P. W. Longest. Validating CFD predictions of highly localized aerosol deposition in airway models: In vitro data and effects of surface properties. *Journal of Aerosol Science*, 59:6–21, 2013.
- P. Honkoop, O. Usmani, and M. Bonini. The current and future role of technology in respiratory care. *Pulmonary Therapy*, pages 1–13, 2022.
- J. Ibanez and J. Raurich. Normal values of functional residual capacity in the sitting and supine positions. *Intensive care medicine*, 8(4):173–177, 1982.

- B. Irving, P. Goussard, R. Gie, A. Todd-Pokropek, and P. Taylor. Segmentation of obstructed airway branches in ct using airway topology and statistical shape analysis. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 447–451. IEEE, 2011.
- B. Irving, T. Douglas, and P. Taylor. 2D X-ray airway tree segmentation by 3D deformable model projection and registration. In *Proceedings of the Fifth International Workshop on Pulmonary Image Analysis*, pages 43–52, 2013.
- M. Ismail, V. Gravemeier, A. Comerford, and W. Wall. A stable approach for coupling multidimensional cardiovascular and pulmonary networks based on a novel pressure-flow rate or pressure-only neumann boundary condition formulation. *International Journal for Numerical Methods in Biomedical Engineering*, 30(4):447–469, 2014.
- A. D. Jagtap, E. Kharazmi, and G. E. Karniadakis. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020.
- A. G.-U. Juarez, R. Selvan, Z. Saghir, and M. de Bruijne. A joint 3D UNet-graph neural network-based method for airway segmentation from chest CTs. In *International Workshop on Machine Learning in Medical Imaging*, pages 583–591. Springer, 2019.
- G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- N. Khajeh-Hosseini-Dalasm and P. W. Longest. Deposition of particles in the alveolar airways: inhalation and breath-hold with pharmaceutical aerosols. *Journal of aerosol science*, 79:15–30, 2015.
- C. Kleinstreuer and Z. Zhang. Airflow and particle transport in the human respiratory system. *Annual review of fluid mechanics*, 42:301–334, 2010.
- T. Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990.
- T. Kohonen. Essentials of the self-organizing map. *Neural networks*, 37:52–65, 2013.

- P. Koullapis, S. C. Kassinos, M. P. Bivolarova, and A. K. Melikov. Particle deposition in a realistic geometry of the human conducting airways: Effects of inlet velocity profile, inhalation flowrate and electrostatic charge. *Journal of biomechanics*, 49(11):2201–2212, 2016.
- P. Koullapis, P. Hofemeier, J. Sznitman, and S. C. Kassinos. An efficient computational fluid-particle dynamics method to predict deposition in a simplified approximation of the deep lung. *European Journal of Pharmaceutical Sciences*, 113:132–144, 2018a.
- P. Koullapis, S. C. Kassinos, J. Muela, C. Perez-Segarra, J. Rigola, O. Lehmkuhl, Y. Cui, M. Sommerfeld, J. Elcner, M. Jicha, et al. Regional aerosol deposition in the human airways: The siminhale benchmark case and a critical assessment of in silico methods. *European Journal of Pharmaceutical Sciences*, 113:77–94, 2018b.
- P. Koullapis, B. Olsson, S. C. Kassinos, and J. Sznitman. Multiscale in silico lung modeling strategies for aerosol inhalation therapy and drug delivery. *Current Opinion in Biomedical Engineering*, 2019.
- P. Koullapis, F. Stylianou, J. Sznitman, B. Olsson, and S. Kassinos. Towards whole-lung simulations of aerosol deposition: A model of the deep lung. *Journal of Aerosol Science*, 144:105541, 2020.
- D. G. Krige. A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139, 1951.
- A. Kuprat, M. Jalali, T. Jan, R. Corley, B. Asgharian, O. Price, R. Singh, S. Colby, and C. Darquenne. Efficient bi-directional coupling of 3D computational fluid-particle dynamics and 1D Multiple Path Particle Dosimetry lung models for multiscale modeling of aerosol dosimetry. *Journal of Aerosol Science*, page 105647, 2020.
- Y. LeCun, Y. Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.
- T.-C. Lee, R. L. Kashyap, and C.-N. Chu. Building skeleton models via 3-D medial surface axis thinning algorithms. *CVGIP: Graphical Models and Image Processing*, 56(6):462–478, 1994.

- X. Li, S. Luo, Q. Hu, J. Li, D. Wang, and F. Chiong. Automatic lung field segmentation in x-ray radiographs using statistical shape and appearance models. *Journal of Medical Imaging and Health Informatics*, 6(2):338–348, 2016.
- J. Liang, Y.-X. Tang, Y.-B. Tang, J. Xiao, and R. M. Summers. Bone suppression on chest radiographs with adversarial learning. In *Medical Imaging 2020: Computer-Aided Diagnosis*, volume 11314, pages 45–50. SPIE, 2020.
- J. Lieman-Sifry, M. Le, F. Lau, S. Sall, and D. Golden. Fastventricle: cardiac segmentation with enet. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 127–138. Springer, 2017.
- K. Lim, D. Harris, P. Sail, and M. Parry. Robust method to predict inspiratory flowrate from the acoustic signature of swirl-based dpis using deep learning. In *Journal of Aerosol Medicine and Pulmonary Drug Delivery*, volume 32, pages A27–A27. Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA, 2019.
- T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- X. Liu, W. H. Doub, and C. Guo. Evaluation of metered dose inhaler spray velocities using phase doppler anemometry (PDA). *International journal of pharmaceutics*, 423(2):235–239, 2012.
- P. Lo, B. Van Ginneken, J. M. Reinhardt, T. Yavarna, P. A. De Jong, B. Irving, C. Fetita, M. Ortner, R. Pinho, J. Sijbers, et al. Extraction of airways from CT (EXACT’09). *IEEE Transactions on Medical Imaging*, 31(11):2093–2107, 2012.
- P. W. Longest, S. Vinchurkar, and T. Martonen. Transport and deposition of respiratory aerosols in models of childhood asthma. *Journal of Aerosol Science*, 37(10):1234–1257, 2006.
- P. W. Longest, G. Tian, R. L. Walenga, and M. Hindle. Comparing mdi and dpi aerosol deposition using in vitro experiments and a new stochastic individual path (sip) model of the conducting airways. *Pharmaceutical research*, 29(6):1670–1688, 2012.

- X. Luo and T. Pedley. A numerical simulation of unsteady flow in a two-dimensional collapsible channel. *Journal of Fluid Mechanics*, 314:191–225, 1996.
- C. MacDonald, R. Hellmuth, L. Priba, E. Murphy, S. Gandy, S. Matthew, R. Ross, and J. Houston. Experimental assessment of two non-contrast mri sequences used for computational fluid dynamics: Investigation of consistency between techniques. *Cardiovascular Engineering and Technology*, 11(4):416–430, 2020.
- S. Marsland, J. Shapiro, and U. Nehmzow. A self-organising network that grows when required. *Neural networks*, 15(8-9):1041–1058, 2002.
- L. Meunier, H. Rakotoarison, P. K. Wong, B. Roziere, J. Rapin, O. Teytaud, A. Moreau, and C. Doerr. Black-box optimization revisited: Improving algorithm selection wizards through massive benchmarking. *IEEE Transactions on Evolutionary Computation*, 2021.
- S. Montesantos, I. Katz, J. Fleming, C. Majoral, M. Pichelin, C. Dubau, B. Piednoir, J. Conway, J. Texereau, and G. Caillibotte. Airway morphology from high resolution computed tomography in healthy subjects and patients with moderate persistent asthma. *The anatomical record*, 296(6):852–866, 2013.
- S. Montesantos, I. Katz, M. Pichelin, and G. Caillibotte. The creation and statistical evaluation of a deterministic model of the human bronchial tree from HRCT images. *PLOS one*, 11(12):e0168026, 2016.
- X. Morales Ferez, J. Mill, K. A. Juhl, C. Acebes, X. Iriart, B. Legghe, H. Cochet, O. De Backer, R. R. Paulsen, and O. Camara. Deep learning framework for real-time estimation of in-silico thrombotic risk indices in the left atrial appendage. *Frontiers in Physiology*, 12:694945, 2021.
- S. Munro, W. Vos, R. Surujbally, and B. Mignot. Using a combination of in-silico lung deposition modelling and secondary endpoint data taken from a phase 3 clinical trial in patients with severe asthma, to provide evidence that lung deposition may be enhanced when using a smart nebuliser instead of a conventional nebuliser. In *Journal of Aerosol Medicine and Pulmonary Drug Delivery*, volume 33, pages A22–A22. Mary Ann Liebert, 2020.

- P. Nardelli, K. A. Khan, A. Corvò, N. Moore, M. J. Murphy, M. Twomey, O. J. O'Connor, M. P. Kennedy, R. S. J. Estépar, M. M. Maher, et al. Optimizing parameters of an open-source airway segmentation algorithm using different CT images. *Biomedical engineering online*, 14(1):62, 2015.
- S. Nousias, E. I. Zacharaki, and K. Moustakas. Avatree: An open-source computational modelling framework modelling anatomically valid airway tree conformations. *PloS one*, 15(4):e0230259, 2020.
- N. Nowak, P. P. Kakade, and A. V. Annapragada. Computational fluid dynamics simulation of airflow and aerosol deposition in human lungs. *Annals of biomedical engineering*, 31(4):374–390, 2003.
- J. M. Oakes, M. Scadeng, E. C. Breen, G. K. Prisk, and C. Darquenne. Regional distribution of aerosol deposition in rat lungs using magnetic resonance imaging. *Annals of biomedical engineering*, 41(5):967–978, 2013.
- J. M. Oakes, A. L. Marsden, C. Grandmont, C. Darquenne, and I. E. Vignon-Clementel. Distribution of aerosolized particles in healthy and emphysematous rat lungs: Comparison between experimental and numerical studies. *Journal of biomechanics*, 48(6):1147–1157, 2015.
- J. M. Oakes, S. C. Shadden, C. Grandmont, and I. E. Vignon-Clementel. Aerosol transport throughout inspiration and expiration in the pulmonary airways. *International journal for numerical methods in biomedical engineering*, 33(9):e2847, 2017.
- J. M. Oakes, S. C. Roth, and S. C. Shadden. Airflow simulations in infant, child, and adult pulmonary conducting airways. *Annals of biomedical engineering*, 46(3):498–512, 2018.
- M. Osanlouy, A. R. Clark, H. Kumar, C. King, M. L. Wilsher, D. G. Milne, K. Whyte, E. A. Hoffman, and M. H. Tawhai. Lung and fissure shape is associated with age in healthy never-smoking adults aged 20–90 years. *Scientific Reports*, 10(1):1–13, 2020.
- J. M. Palmer, J. W. Bok, S. Lee, T. R. Dagenais, D. R. Andes, D. P. Kontoyiannis, and N. P. Keller. Loss of ccla, required for histone 3 lysine 4 methylation, decreases growth but increases secondary metabolite production in *aspergillus fumigatus*. *PeerJ*, 1:e4, 2013.

- A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*, 2016.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- P. Perdikaris and G. E. Karniadakis. Model inversion via multi-fidelity bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. *Journal of The Royal Society Interface*, 13(118):20151107, 2016.
- K. Poorbahrami and J. M. Oakes. Regional flow and deposition variability in adult female lungs: A numerical simulation pilot study. *Clinical Biomechanics*, 66:40–49, 2019.
- N. Pozin, S. Montesantos, I. Katz, M. Pichelin, I. Vignon-Clementel, and C. Grandmont. A tree-parenchyma coupled model for lung ventilation simulation. *International journal for numerical methods in biomedical engineering*, 33(11):e2873, 2017.
- M. Raissi, A. Yazdani, and G. E. Karniadakis. Hidden fluid mechanics: A Navier-Stokes informed deep learning framework for assimilating flow visualization data. *arXiv preprint arXiv:1808.04327*, 2018.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- J. Rapin and O. Teytaud. Nevergrad - A gradient-free optimization platform. <https://GitHub.com/FacebookResearch/Nevergrad>, 2018.
- S. M. Rausch, D. Haberthür, M. Stampanoni, J. Schittny, and W. Wall. Local strain distribution in real three-dimensional alveolar geometries. *Annals of biomedical engineering*, 39(11):2835–2843, 2011.

- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- C. J. Roth, L. Yoshihara, M. Ismail, and W. A. Wall. Computational modelling of the respiratory system: discussion of coupled modelling approaches and two recent extensions. *Computer Methods in Applied Mechanics and Engineering*, 314:473–493, 2017.
- J. Russo, R. Robinson, and M. J. Oldham. Effects of cartilage rings on airflow and particle deposition in the trachea and main bronchi. *Medical engineering & physics*, 30(5):581–589, 2008.
- N. Sarkalkan, H. Weinans, and A. A. Zadpoor. Statistical shape and appearance models of bones. *Bone*, 60:129–140, 2014.
- A. Scheinherr, L. Bailly, O. Boiron, A. Lagier, T. Legou, M. Pichelin, G. Caillibotte, and A. Giovanni. Realistic glottal motion and airflow rate during human breathing. *Medical engineering & physics*, 37(9):829–839, 2015.
- A. A. A. Setio, A. Traverso, T. De Bel, M. S. Berens, C. Van Den Bogaard, P. Cerello, H. Chen, Q. Dou, M. E. Fantacci, B. Geurts, et al. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Medical image analysis*, 42:1–13, 2017.
- L. Shen, W. Zhao, and L. Xing. Patient-specific reconstruction of volumetric computed tomography images from a single projection view via deep learning. *Nature biomedical engineering*, 3(11):880–888, 2019.
- M. Sonka, W. Park, and E. A. Hoffman. Rule-based detection of intrathoracic airway trees. *IEEE transactions on medical imaging*, 15(3):314–326, 1996.
- B. Su, J.-M. Zhang, H. Zou, D. Ghista, T. T. Le, and C. Chin. Generating wall shear stress for coronary artery in real-time using neural networks: Feasibility and initial results based on idealized models. *Computers in biology and medicine*, 126:104038, 2020.

- J. Sznitman. Respiratory microflows in the pulmonary acinus. *Journal of biomechanics*, 46(2): 284–298, 2013.
- H. Tang, C. Zhang, and X. Xie. Automatic pulmonary lobe segmentation using deep learning. *arXiv preprint arXiv:1903.09879*, 2019.
- M. H. Tawhai, P. Hunter, J. Tschirren, J. Reinhardt, G. McLennan, and E. A. Hoffman. Ct-based geometry analysis and finite element models of the human and ovine bronchial tree. *Journal of applied physiology*, 97(6):2310–2321, 2004.
- The Global Asthma Network. The Global Asthma Report, 2018.
- G. Tian, M. Hindle, S. Lee, and P. W. Longest. Validating CFD predictions of pharmaceutical aerosol deposition with in vivo data. *Pharmaceutical research*, 32(10):3170–3187, 2015.
- J. Tschirren, E. A. Hoffman, G. McLennan, and M. Sonka. Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans. *IEEE transactions on medical imaging*, 24(12):1529–1539, 2005a.
- J. Tschirren, G. McLennan, K. Palágyi, E. A. Hoffman, and M. Sonka. Matching and anatomical labeling of human airway tree. *IEEE transactions on medical imaging*, 24(12):1540–1547, 2005b.
- J. Tschirren, T. Yavarna, and J. Reinhardt. Airway segmentation framework for clinical environments. In *Proceedings of 2nd international workshop on pulmonary image analysis*, pages 227–38, 2009.
- S. P. Väänänen, L. Grassi, G. Flivik, J. S. Jurvelin, and H. Isaksson. Generation of 3D shape, density, cortical thickness and finite element mesh of proximal femur from a DXA image. *Medical image analysis*, 24(1):125–134, 2015.
- J. F. van Boven, B. Cushen, I. Sulaiman, G. Greene, E. MacHale, M. C. Mokoka, F. Doyle, R. B. Reilly, K. Bennett, and R. W. Costello. Personalising adherence-enhancing interventions using a smart inhaler in patients with COPD: an exploratory cost-effectiveness analysis. *NPJ primary care respiratory medicine*, 28(1):1–3, 2018.

- H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Computers in physics*, 12(6):620–631, 1998.
- J. Williams, J. Kolehmainen, S. Cunningham, A. Ozel, and U. Wolfram. Effect of patient inhalation profile and airway structure on drug deposition in image-based models with particle-particle interactions. *International Journal of Pharmaceutics*, 612:121321, 2022a.
- J. Williams, A. Ozel, and U. Wolfram. jvwilliams23/deepLungFoam: v0.1.0, 2022b. URL <https://doi.org/10.5281/zenodo.7448214>.
- J. Williams, U. Wolfram, and A. Ozel. Neural stochastic differential equations for particle dispersion in large-eddy simulations of homogeneous isotropic turbulence. *Physics of Fluids*, 34(11):113315, 2022c.
- J. Williams, A. Ozel, and U. Wolfram. pyssam – a python library for statistical modelling of biomedical shape and appearance. *arXiv preprint arXiv:2301.04416*, 2023.
- J. Xi, J. E. Yuan, M. Yang, X. Si, Y. Zhou, and Y.-S. Cheng. Parametric study on mouth–throat geometrical factors on deposition of orally inhaled aerosols. *Journal of Aerosol Science*, 99: 94–106, 2016.
- W. Yang, Y. Chen, Y. Liu, L. Zhong, G. Qin, Z. Lu, Q. Feng, and W. Chen. Cascade of multi-scale convolutional neural networks for bone suppression of chest radiographs in gradient domain. *Medical image analysis*, 35:421–433, 2017.
- M. Yin, X. Zheng, J. D. Humphrey, and G. E. Karniadakis. Non-invasive inference of thrombus material properties with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 375:113603, 2021.
- K. Youn, M. S. Park, and J. Lee. Iterative approach for 3D reconstruction of the femur from un-calibrated 2D radiographic images. *Medical engineering & physics*, 50:89–95, 2017.
- J. Yun, J. Park, D. Yu, J. Yi, M. Lee, H. J. Park, J.-G. Lee, J. B. Seo, and N. Kim. Improvement of fully automated airway segmentation on volumetric computed tomographic images using a 2.5 dimensional convolutional neural net. *Medical image analysis*, 51:13–20, 2019.

- Y. Zhang and W. H. Finlay. Measurement of the effect of cartilaginous rings on particle deposition in a proximal lung bifurcation model. *Aerosol Science and Technology*, 39(5): 394–399, 2005.
- J. Zhao, Y. Feng, and C. A. Fromen. Glottis motion effects on the particle transport and deposition in a subject-specific mouth-to-trachea model: A cfpd study. *Computers in Biology and Medicine*, 116:103532, 2020.
- G. Zheng, L.-P. Nolte, and S. J. Ferguson. Scaled, patient-specific 3D vertebral model reconstruction based on 2D lateral fluoroscopy. *International journal of computer assisted radiology and surgery*, 6(3):351–366, 2011.

Supplementary material

Josh Williams¹, Haavard Ahlqvist¹, Alexander Cunningham¹, Andrew Kirby², Ira Katz³, John Fleming^{4,5}, Joy Conway^{4,6}, Steve Cunningham⁷, Ali Ozel^{1*}, and Uwe Wolfram^{1*}

¹School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK

²Royal Hospital for Sick Children, NHS Lothian, Edinburgh, UK

³Consultant, Meudon, France

⁴National Institute of Health Research Biomedical Research Centre in Respiratory Disease, Southampton, UK

⁵Department of Medical Physics and Bioengineering, University Hospital Southampton NHS Foundation Trust, Southampton, UK

⁶Respiratory Sciences, Centre for Health and Life Sciences, Brunel University, London, UK

⁷Centre for Inflammation Research, University of Edinburgh, Edinburgh, UK

March 3, 2023

* UW and AO share last authorship.

Address correspondence to u.wolfram@hw.ac.uk or a.ozel@hw.ac.uk

1 Manual cleanup of upper airways

To mitigate the issues in image acquisition discussed by Conway et al. (2012), we had to manually clean the throat to allow passage for the air and particles to travel through which would be close to expected *in vivo* conditions during inhalation. We extracted a baseline upper airway segmentation using a semi-automatic region-growing approach implemented in Python. The baseline segmentation was imported into 3D Slicer, where the glottal cross-sectional area was filled in manually with a paintbrush tool to better agree with upper airway morphologies in literature (Feng et al., 2018; Zhao et al., 2020; Scheinherr et al., 2015). The segmentation was then smoothed to remove artificial bumps introduced by the paintbrush tool. Visual representations of two corrected airways are shown in Figure S1.

2 CNN architecture analysis

To find the optimal loss function γ parameter in Equation 11, we trained three U-Net CNNs with $\gamma = \{1, 3, 5\}$ (Figure S2). By comparing the DICE coefficient over the 10 cases in the validation set, we observed $\gamma = 1$ to produce the lowest minimum DICE coefficient (0.85). When $\gamma = 3$ and $\gamma = 5$, the minimum and maximum DICE coefficients were approximately the same (0.855 minimum and 0.95 maximum). Throughout the remainder of the study, we chose $\gamma = 5$ in our loss function due to the slightly higher lower quartile, median and upper quartile (all approximately 1% larger for $\gamma = 5$).

As the U-Net has a large memory consumption, we also trained an ENet CNN architecture as a less computationally expensive alternative (Figure S3a). Both architectures were trained using $\gamma = 5$ and using 25 augmentations (with rotation $\pm 15^\circ$). There was a significant difference in DICE coefficient ($p < 0.05$), as the U-Net median and upper-quartile was 4% larger than the ENet. The U-Net lower-quartile was 4.7% larger than the E-Net. We also compared the time-taken for one forward-pass through the networks (inference time), for all images in the training and validation dataset (Figure S3b). For the smallest images in the dataset (number of voxels below the 25th percentile), the mean inference time was 5.89 s and 7.81 s for the ENet and U-Net, respectively (1.3 times speed-up with ENet). In contrast, for the largest images in the dataset (number of voxels larger than the 75th percentile), the

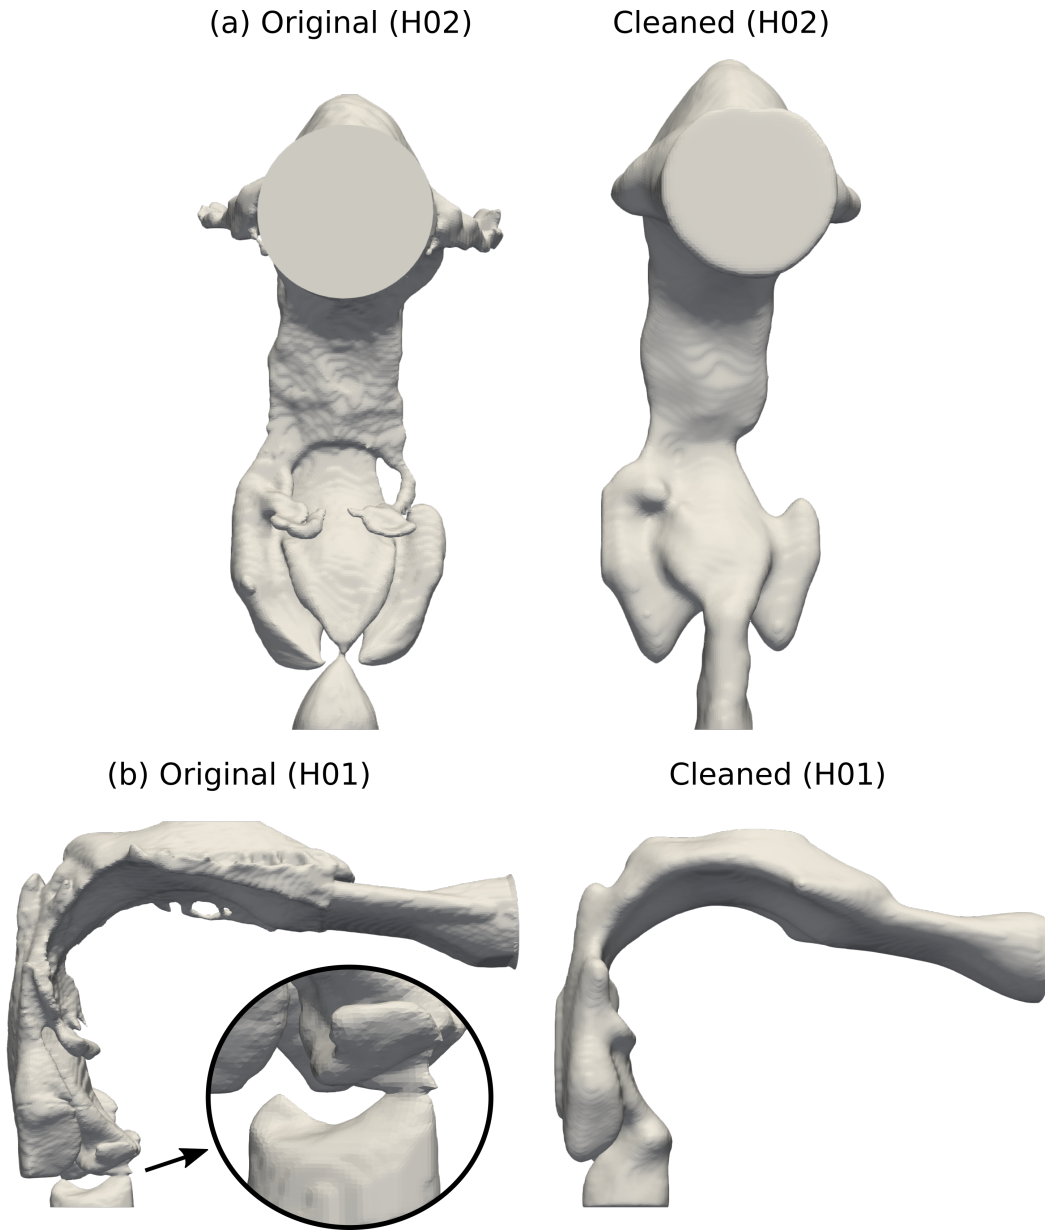


Figure S1: Visualisation of upper airway surface mesh before and after manual cleanup for two cases (Conway et al., 2012). The upper row shows the most constricted case (H02). The lower row shows a less extreme, but still heavily constricted case with a zoomed in view.




Figure S2: Sensitivity analysis of focal loss parameter γ in Equation 11 to optimise CNN DICE coefficient. CNNs were trained with no augmentations and results were computed from 10 scans in validation set.

mean inference time was 15.9s and 91.28s for the ENet and U-Net, respectively, which is a speed-up of 5.7 times when using the ENet. As speed is not a significant driver of our developments at the present stage, we use the U-Net as our default segmentation tool for this study due to its improved DICE coefficient. However, due to memory overheads required for inference with the U-Net, we used the ENet to segment airways from the high resolution CT of the Southampton/ Air Liquide dataset. Additionally, the ENet DICE coefficient is reasonable and may be preferable for implementation in a clinical workflow, where speed is a key factor.

3 SSAM morphology assessment

Here we provide the results presented in Figures 3 and 5 as scatter plots to allow for comparison of absolute values of lung space volume and airway diameter. We found it important to provide both plots, to show that the absolute values are within a realistic range. This also shows a clear linear correlation between the ground truth measurements and those predicted by the SSAM. The lung space volume predicted by the SSAM agreed particularly well (concordance correlation coefficient $CCC \approx 0.89$, Figure S4). The diameter predicted by the SSAM (Figure S5) showed a less ideal fit than the lung volume, as $CCC = 0.657$ and 0.71 . The fit was slightly improved with two projections, compared to one projection (5.3%




Figure S3: Comparison of accuracy and speed for U-Net and ENet architectures. Panels show (a) DICE coefficient from 10 scans in validation set, and (b) inference time relative to size of image, computed on combined training and validation set.

increase).

4 Choice of generated airway diameter model

As the SSAM-reconstructed airways do not have the same level of information as the CT-based airways, best practices for generating distal airways may not be the same as that of airways segmented directly from patient CT data. Specifically, on the chest X-ray, only the trachea and main bronchi are visible. Even for these ‘visible’ airways, the contrast is poor and the quality of reconstructed diameters was found to be varied in our analysis. When using the diameter of the parent branch to assign diameters of the entire airway tree (Bordas et al., 2015), it makes sense that erroneous diameters cause some level of error that propagates into the distal airways. Therefore, we compared various models for computing diameter obtained in literature (Bordas et al., 2015; Montesantos et al., 2016).

Bordas et al. (2015) assigned a logarithmic decay of diameter based on Horsfield order (Horsfield et al., 1976) defined as

$$\log D_i(H_i) = (H_i - N_H) \log(R_d H) + \log(D_N) \quad (1)$$




Figure S4: Comparison of lung volume predicted by our SSAM compared to the ground truth segmentations. We show the influence of including an additional projection on lung volume error. Panels show SSAM results with (a) one DRR provided for fitting (anterior-posterior projection), and (b) two DRRs provided (anterior-posterior and lateral projections). CCC is the concordance correlation coefficient.




Figure S5: Comparison of airway diameter predicted by our SSAM compared to the ground truth segmentations. We show the influence of including an additional projection on diameter error in the trachea and main bronchi. Panels show SSAM results with (a) one DRR provided for fitting (anterior-posterior projection), and (b) two DRRs provided (anterior-posterior and lateral projections). CCC is the concordance correlation coefficient.

where D_i is the diameter of branch i , H is the Horsfield order, N_H is the Horsfield order of the reference airway in the image-based domain, D_N is the diameter of the reference airway in the image-based domain, $R_d H = 1.15$ is a constant representing the logarithmic decrease of diameter with Horsfield order (Horsfield et al., 1976; Tawhai et al., 2004; Bordas et al., 2015). This approach requires declaration of an airway segment in the image-based airways to obtain reference values that are used to assign diameter and Horsfield order to the generated airways. Bordas et al. (2015) used the parent branch of the image-based airway that is connected to the generated airway, i (referred to here as ‘parent’ model). Alternatively, Montesantos et al. (2016) calculated airway diameter based on the airway length as $D_i = L_i/3$, with a cutoff $D_{child} \leq 0.95D_{parent}$ (referred to here as ‘length’ model). Ventilation simulations using the ‘parent’ model was shown to have good agreement with clinical ventilation data in healthy and diseased patients (Bordas et al., 2015). Therefore, we used the parent model for CT-based airways. We used ground truth segmentations with the parent model as a benchmark to evaluate the accuracy of the parent model or length model in on generated airway diameter in SSAM-reconstructed airways.

We found the length model to produce better agreement with the ground truth dataset than the parent model proposed by Bordas et al. (2015) (Figure S6a). This is likely due to difficulty inferring the diameter of the central airways from a chest X-ray image. As can be seen in Figure S6b, the generated airway length showed excellent agreement in the SSAM and ground truth datasets. Based on this, we chose to use the length model of Montesantos et al. (2016) to assign diameters to the full SSAM airway tree.

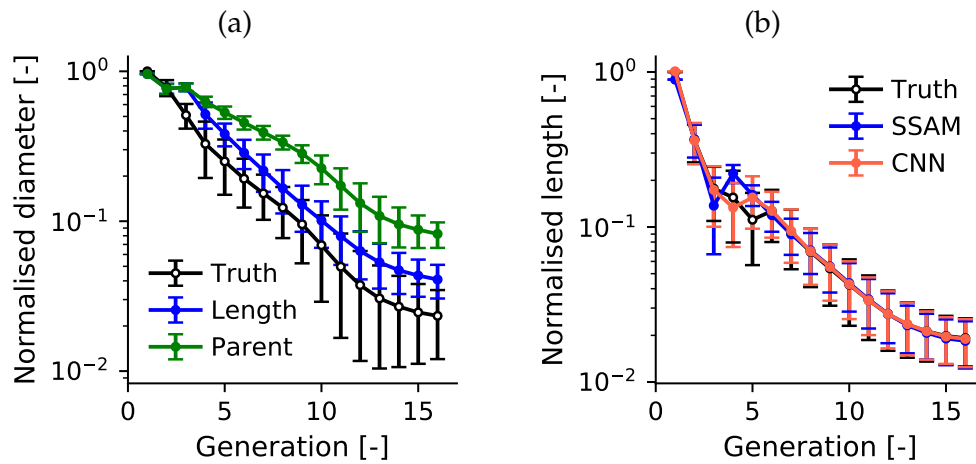


Figure S6: Comparison of diameter error in full conducting airway tree from a SSAM with varying models to calculate diameter. The ‘truth’ results were generated from ground truth segmentations with the ‘parent’ model for diameter (Bordas et al., 2015).

References

- R. Bordas, C. Lefevre, B. Veeckmans, J. Pitt-Francis, C. Fetita, C. E. Brightling, D. Kay, S. Siddiqui, and K. S. Burrowes. Development and analysis of patient-based complete conducting airways models. *PloS one*, 10(12):e0144105, 2015.
- J. Conway, J. Fleming, C. Majoral, I. Katz, D. Perchet, C. Peebles, L. Tossici-Bolt, L. Collier, G. Caillibotte, M. Pichelin, et al. Controlled, parametric, individualized, 2-D and 3-D imaging measurements of aerosol deposition in the respiratory tract of healthy human subjects for model validation. *Journal of aerosol science*, 52:1–17, 2012.
- Y. Feng, J. Zhao, C. Kleinstreuer, Q. Wang, J. Wang, D. H. Wu, and J. Lin. An in silico inter-subject variability study of extra-thoracic morphology effects on inhaled particle transport and deposition. *Journal of Aerosol Science*, 123:185–207, 2018.
- K. Horsfield, F. G. Relea, and G. Gunning. Diameter, length and branching ratios in the bronchial tree. *Respiration physiology*, 26(3):351–356, 1976.
- S. Montesantos, I. Katz, M. Pichelin, and G. Caillibotte. The creation and statistical evaluation of a deterministic model of the human bronchial tree from HRCT images. *PLOS one*, 11(12):e0168026, 2016.
- A. Scheinherr, L. Bailly, O. Boiron, A. Lagier, T. Legou, M. Pichelin, G. Caillibotte, and A. Gio-

- vanni. Realistic glottal motion and airflow rate during human breathing. *Medical engineering & physics*, 37(9):829–839, 2015.
- M. H. Tawhai, P. Hunter, J. Tschirren, J. Reinhardt, G. McLennan, and E. A. Hoffman. Ct-based geometry analysis and finite element models of the human and ovine bronchial tree. *Journal of applied physiology*, 97(6):2310–2321, 2004.
- J. Zhao, Y. Feng, and C. A. Fromen. Glottis motion effects on the particle transport and deposition in a subject-specific mouth-to-trachea model: A cfpd study. *Computers in Biology and Medicine*, 116:103532, 2020.