

Modelling Airway Geometry as Stock Market Data using Bayesian Change-point Detection

Abstract. Numerous lung diseases, such as idiopathic pulmonary fibrosis (IPF), are manifested as dilation of the airways. In order to quantify the progression of such diseases, prior analyses have introduced a variety of computational methods for measuring the airway cross-sectional area from CT images. However, the combination of image noise and anatomical fluctuations causes high variability in the measurements of cross-sectional areas, rendering the identification of affected regions very difficult. In this work, we introduce a noise-robust method for automatically detecting the location of progressive airway dilatation given two measurements of the same airway acquired at different time points. To this end, we propose a probabilistic model of abrupt relative variations between two measurements of cross-sectional areas and perform inference via Reversible Jump Markov Chain Monte Carlo sampling. We demonstrate the efficacy of the proposed method on two datasets; (i) images of healthy airways with simulated dilatation; (ii) pairs of real images of IPF-affected airways acquired at 1 year intervals. Results on simulated data show that our model is able to detect the starting location of airway dilatation within the accuracy of 2.5mm, while the experiments on the IPF dataset display reasonable agreement with radiologists. In addition, we demonstrate that by estimating the location of dilatation we can compute a relative change in airway volume which is useful for quantification of IPF disease progression.

1 Introduction

In fibrosing lung disease, contraction of the lung interstitium pulls on airway walls, and causes a dilatation. Airway dilatation is calculated using crude visual scores have been shown to be powerful predictors of outcome in idiopathic pulmonary fibrosis (IPF) [6]. Computed Tomography (CT) coupled with image analysis algorithms to measure the airway lumen can provide an excellent tool to quantitatively assess these dilatation. CT gives a very high image contrast between air within the airway lumen and the surrounding airway wall. However, airway lumen measuring algorithms are subjected to sources of noise. Variations in voxel sizes and CT reconstruction algorithms have been shown to affect the precision of the cross-sectional area measurements of the airways [5]. In addition, there can be variations in airway cross-sectional area due to normal biological processes such as points of airway bifurcation where the airway transiently dilates [11, 2]. These sources of noise makes post-processing of measurements difficult to locate dilatation on an airway track. Identifying the point of dilatation and subsequently calculating the volume of abnormal airways in the lung potentially

provides a sensitive means for quantifying subtle worsening of disease in IPF across longitudinal scans.

In this work, we introduce a variant of Bayesian Changepoint Detection (BCPD) models to automatically identify the location of airway dilatation in the presence of strong measurement noise, given two measurements of the same airway acquired at different time points. BCPD model, typically used in DNA sequencing [12] and stock market data analysis [9], aims to capture abrupt variations in the underlying distributions of a given signal or a time series. The method processes a series of airway cross-sectional area changes between baseline (first) and follow-up (second) CT scans, and generates the posterior distribution over multiple possible points of abnormal variations, whose the mode is taken as the final prediction. We test the efficacy of the method on (1) CT images of healthy airways with simulated dilatation, and (2) pairs of real images of IPF-affected airways acquired approximately 1 year intervals. For the simulated dataset, we measured the detection accuracy with respect to the commonly used naive thresholding and maximum likelihood based methods [8]. For the longitudinal IPF dataset, we compare the predictions of our model to measurements from radiologists based on two different protocols.

2 Method

In this section, we introduce our method for quantifying progression of IPF on a patients serial CT scans. The method proceeds as follows. Firstly, we fit a tubular shape model to the airways in both baseline and follow-up CT scans, and acquire estimates of the cross-sectional areas along them (Sec. 2.1). We then treat the difference in the cross-sectional areas as a 1D signal, and employ the proposed Bayesian changepoint model to estimate the posterior distribution over locations of abrupt airway dilatation (Sec. 2.2 and 2.3). Lastly, we post-process this posterior distribution to determine the region of dilatation (Sec. 2.4).

2.1 Airway Pre-processing

In this work, for each airway track, we acquire a series of cross-sectional area measurements using the method proposed by Quan et al [11] (Fig. 1). Following a semi-manual segmentation of the airway, the method computes the airway centreline and the corresponding normal planes at contiguous intervals, each of which is then used to estimate the cross-sectional area. The final output is a 1D function of cross-sectional area along the airways for baseline $f_B(x)$ and the follow-up $f_F(x)$ scans.

The next task was to align the signals on both baseline $f_B(x)$ and follow-up $f_F(x)$ scans. To this end, we resample the signal to 1mm using cubic interpolation. We considered the first 50 points on both signals g_B, g_F from the start of the carina. To register the two airways together, we apply the transformation $f_F(x - a)$ where $a = \operatorname{argmin}_{a \in [-5, 5]} \left\| \log \left(\frac{g_B(x)}{g_F(x-a)} \right) \right\|_2$. The longest of the two signal were truncated from the right hand side such that both signals were of




Fig. 1. Summary of the pipeline developed by Quan et al [11]. We have implemented the method as part of the pre-processing stage to model the geometry as a 1D signal.

the same length. For the rest of the methodology, we will only consider the series difference defined as $\mathbf{y} = \log(f_F) - \log(f_B)$.

2.2 Bayesian Changepoint Model

The progression of fibrosing lung diseases can manifest itself as dilatation of the internal airways. Thus, we hypothesise that at the start of dilatation, the series \mathbf{y} undergoes an abrupt variation, which we refer to as a changepoint. More formally, given signal $\mathbf{y} = (y_1, \dots, y_n)$ of length n , we define a changepoint τ as the location where there exists a change in parameters θ in the underlying distribution F . In other words, at changepoint $1 < \tau < n$, the observations \mathbf{y} can be separated at τ such that:

$$\mathbf{y} = \begin{cases} (y_1, \dots, y_\tau) \sim F(\theta_1) \\ (y_{\tau+1}, \dots, y_n) \sim F(\theta_2) \end{cases}, \quad (1)$$

where $\theta_1 \neq \theta_2$. This definition can be naturally extended to the scenario with M changepoints; we denote the changepoint location vector by $\tau = (\tau_1, \dots, \tau_k)$, with parameters $\theta = (\theta_1, \dots, \theta_{k+1})$ for each respective segment. For ease of notation, we also denote $\tau_0 = 1$ and $\tau_{k+1} = n$. Assuming statistical independence between segments, the likelihood factorises as: $p(\mathbf{y}|\tau, \theta, M) = \prod_{l=1}^{k+1} F(y_{\tau_{l-1}:\tau_l}|\theta_l)$, where $y_{\tau_{l-1}:\tau_l} = (y_{\tau_{l-1}}, \dots, y_{\tau_l})$. We also specify prior distributions on the the number of changepoints $p(M; \delta)$, the locations of the changepoints $p(\tau|M; \gamma)$, and the parameters of the corresponding segments $p(\theta|M; \beta)$ where β , γ and δ represent the hyper-parameters.

Given the likelihood and the prior distributions above, we would like to estimate the posterior distributions over the number of changepoints $p(M|\mathbf{y})$, the locations of changepoints $p(\tau|\mathbf{y})$ and the parameters of the respective segments $p(\theta|\mathbf{y})$. Commonly in Bayesian changepoint analysis, likelihoods with a conjugate prior are chosen to allow calculation of analytical posteriors for changepoints and parameters, enabling faster convergence and better accuracy [3]. However, in our work, we wish to relax this conjugacy assumption in order to allow flexibility in the choices of prior and likelihoods. In the next section, we describe the inference scheme in such highly general setting, viable for arbitrary forms of the likelihood distribution F and the prior. In this paper, F was defined as the Student t distribution with parameters $\theta_l = (\nu_l, \mu_l, \sigma_l)$ with degrees of freedom

ν_l , mean μ_l and standard deviation σ_l . We chose t-distribution for its robustness to outliers [10] caused by noise from the area measurements.

2.3 Reversible Jump MCMC for Posterior Inference

Posterior inference with our model possesses two challenges. Firstly, without the conjugacy assumption, computing the posterior distributions is intractable. Secondly, the dimensionality of the posterior distribution over the changepoints τ is given by M and varies during inference. To combat the first problem, we use the Metropolis-Hasting (MH) algorithm [1], a variant of Markov Chain Monte Carlo (MCMC) methods that can sample from the posterior, with or without conjugacy. Given that the number of changepoints M is known, MH can be used to sample from the posterior distributions over the changepoints τ and segment parameters θ . To address the second problem of varying posterior dimensionality M , we extend the above sampling scheme to the Reversible Jump MCMC framework [4]. Taken all together, the method is capable of traversing the full posterior distributions for M, τ, θ and we refer to this as Reversible Jump Metropolis Hasting (RJMH) algorithm.

The RJMH proceeds by randomly executing one of four possible moves at each iteration: (i) resample parameters θ ; (ii) move an existing changepoint; (iii) add a new changepoint; (iv) delete an existing changepoint (see Algorithm 1 for the pseudo-code). We will denote the probability of these RJMCMC moves as $\underline{\gamma} = \{\gamma_r, \gamma_m, \gamma_a, \gamma_d\}$, where $\gamma_r + \gamma_m + \gamma_a + \gamma_d = 1$. We also define the maximum number of changepoints k_{max} and at the boundary cases for k , we impose restrictions such that $\gamma_d = 0$ for $k = 0$ and $\gamma_a = 0$ for $k = k_{max}$. Each move updates the appropriate subset of parameters θ, τ by sampling from the corresponding proposal distributions $q(\theta_{new}|\theta_{old})$ and $q(\tau_{new}|\tau_{old})$, and is only executed if it passes the associated acceptance criteria α . An overview of these moves can be described as: (i) The parameters of every segment is updated. (ii) One of the changepoints is selected randomly and updated, and the parameters of its neighbouring two segments are updated. (iii) Randomly choose a location in the time series and add it as a changepoint. The parameters of the resultant two new segments are also sampled. (iv) One of the changepoints is randomly removed and the parameters of the new segment are sampled. The proposal distributions and acceptance criteria of these moves take different forms, and next we describe them in detail.

We can subdivide these four moves into two sets: fixed M and changing M i.e. resampling parameters and moving changepoints does not affect the changepoint model M whereas adding or deleting a changepoint changes the value of M . We first begin by describing the fixed M moves: (i) The algorithm resamples parameters of each segment by proposing Gaussian perturbations around the current values of parameters for all segments i.e. $q(\theta_{new}|\theta_{old}) = N(\theta_{old}, \epsilon) = (\mu_{old}, \sigma_{old}^2, \nu_{old}) + N(0, \epsilon)$. (ii) The algorithm selects a changepoint τ at random and shifts it with a Poisson perturbation i.e. $q(\tau_{new}|\tau_{old}) = \tau_{old} + (-1)^b \text{Poi}(\lambda)$ where $b \sim \text{Binary}[0, 1]$. The segments neighbouring this new changepoint location has its parameters θ resampled as in move (i) using the current segment

parameters. Note that each proposal in the above moves will be symmetrical and therefore, cancel out with the reverse proposal term (See Algorithm 1). For the changing M moves, the algorithm moves in dimensionality and we need to introduce alternative proposal terms for the new segments produced as there are no current segment parameters to base our proposals on. The following describes the changing M moves: (iii) The algorithm proposes a new changepoint randomly over the length of the time series (excluding the final point as no changepoint can exist there) i.e. $\tau_{new} \sim U[1, n-1]$ and $M_{new} = M + 1$. The proposed changepoint will split an existing segment into two new segments, which each will require a new set of parameters. Our proposal for μ and σ^2 in each new segment is defined by a Gaussian perturbation on the empirical values of each respective parameter i.e. for the new segment *left* of the new changepoint $q(\mu_l|\mu_{old}) = \bar{y}_l + N(0, \epsilon_1^2)$ and $q(\sigma_l^2|\sigma_{old}^2) = s_l^2 + N(0, \epsilon_2^2)$ where \bar{y}_l and s_l^2 are the empirical mean and variance of the new *left* segment respectively. The new *right* segment follows analogously for μ and σ^2 . For the degrees of freedom ν , we use the proposal on the new left segment $q(\nu_l|\nu_{old}) = \nu_{old} + u$ and $q(\nu_r|\nu_{old}) = \nu_{old} - u$ for the new right segment where $u \sim N(0, \epsilon^2)$. By enforcing a dependence on the current parameters for ν_l and ν_r , we introduce a Jacobian term. Combining the reverse move described in the next step, the Jacobian can be shown to be $|J| = 2$. (iv) The algorithm selects an existing changepoint and proposes to delete it, i.e. $M_{new} = M - 1$, introducing a new single segment to replace the two segments left and right of the changepoint. Our proposal for μ and σ^2 in the new single segment is defined by a Gaussian perturbation on the empirical values of each respective parameter i.e. for the new segment $q(\mu_{new}|\mu_l, \mu_r) = \bar{y}' + N(0, \epsilon^2)$ and $q(\sigma_{new}^2|\sigma_l^2, \sigma_r^2) = \tilde{s}^2 + N(0, \epsilon^2)$ where \bar{y}' and \tilde{s}^2 are the empirical mean and variance of the new single segment respectively. For the degrees of freedom ν , we use the proposal on the new single segment $q(\nu_{new}|\nu_l, \nu_r) = (\nu_l + \nu_r)/2$, which is a deterministic proposal based on the variables generated in the reverse move, therefore only introducing the reciprocal Jacobian term i.e. $|J| = 1/2$. Note that $q(M_{new}|M)$ and $q(M|M_{new})$ in both RJ-steps, will be symmetrical since we restrict the algorithm to add and remove only one changepoint in these moves.

The overall algorithm structure can be seen in the pseudo-code Algorithm 1. With the model defined, we mentioned the priors used and address some common implementation techniques. The priors used are given as follows: $\mu \sim N(0, 1)$, $\sigma^2 \sim \text{Inv-}\chi^2(5, 0.4^2)$, $\nu \sim U[2, 100]$ and $M \sim \text{Bin}(n-1, \frac{0.5}{n-1})$. The hyperparameters for μ, σ^2, ν were chosen to be non-informative and within plausible ranges. In order to reduce the number of changepoints we could detect (in the case of acquisition noise from the CT), we set the expectation for M to be sufficiently low. In terms of implementation, we follow a standard procedure of setting a burn-in for the number of iterations to ignore any potential issues with initialisation at approximately 25% of the total iteration count. To remove any chance of auto-correlation, we thin the number of samples by only storing the 5th iteration, after the burn-in period. Additionally, we choose the number of iterations, N , to be sufficiently large in order to achieve convergence.

Algorithm 1 RJMH Changepoint Detection [4].

Data: Single airway time series data, $\mathbf{y} = (y_1, \dots, y_n)$

Model Parameters: $M = \#$ changepoints, $\theta =$ segment parameters, $\tau =$ changepoint location

RJMCMC Parameters: $N =$ iterations, $B =$ burn-in, $T =$ Thinning

Initialisation: Set $s = (M, \theta, \tau)$

Initialise set s with no changepoints and empirical segment parameters

for N iterations **do**

Sample $u \sim \text{Uniform}[0, 1]$

if $u \leq \Upsilon_r$ **then**

MH-step (1): Resample Segment Parameters

 Resample $\tilde{\theta} \sim q(\tilde{\theta}|\theta)$ for each segment

 Compute acceptance ratio:

$$\alpha(\tilde{s}, s) = \min \left\{ 1, \frac{p(\mathbf{y}|\tau, \tilde{\theta}, M) p(\tilde{\theta}) q(\theta|\tilde{\theta})}{p(\mathbf{y}|\tau, \theta, M) p(\theta) q(\tilde{\theta}|\theta)} \right\}$$

 Generate $v \sim \text{Uniform}[0, 1]$

if $v \geq \alpha$ **then**

 Accept new state

$s = (M, \tilde{\theta}, \tau)$

end if

else if $\Upsilon_r < u \leq \Upsilon_r + \Upsilon_m$ **then**

MH-step (2): Move changepoint

 Move a changepoint $\tilde{\tau} \sim q(\tilde{\tau}|\tau)$ and propose new segment parameters $\tilde{\theta} \sim q(\tilde{\theta}|\theta)$

 Compute acceptance ratio:

$$\alpha(\tilde{s}, s) = \min \left\{ 1, \frac{p(\mathbf{y}|\tilde{\tau}, \tilde{\theta}, M) p(\tilde{\theta}) q(\tau|\tilde{\tau}) q(\theta|\tilde{\theta})}{p(\mathbf{y}|\tau, \theta, M) p(\theta) q(\tilde{\tau}|\tau) q(\tilde{\theta}|\theta)} \right\}$$

 Generate $v \sim \text{Uniform}[0, 1]$

if $v \geq \alpha$ **then**

 Accept new state

$s = (M, \tilde{\theta}, \tilde{\tau})$

end if

else if $\Upsilon_r + \Upsilon_m < u \leq \Upsilon_r + \Upsilon_m + \Upsilon_a$ **then**

RJ-step (3): Add changepoint

 Propose new changepoint $\tilde{\tau}$, propose new segment parameters $\tilde{\theta}$

 Compute acceptance ratio

$$\alpha(\tilde{s}, s) = \min \left\{ 1, \frac{p(\mathbf{y}|\tilde{\tau}, \tilde{\theta}, \tilde{M}) p(\tilde{\tau}) p(\tilde{\theta}) p(\tilde{M}) q(\tau|\tilde{\tau}) q(\theta|\tilde{\theta}) q(M|\tilde{M})}{p(\mathbf{y}|\tau, \theta, M) p(\tau) p(\theta) p(M) q(\tilde{\tau}|\tau) q(\tilde{\theta}|\theta) q(\tilde{M}|M)} |J| \right\}$$

 where $|J| = 2$ and $q(\tilde{M}|M)$ is symmetrical to $q(M|\tilde{M})$.

 Generate $v \sim \text{Uniform}[0, 1]$

if $v \geq \alpha$ **then**

 Accept new state

$s = (M, \tilde{\theta}, \tilde{\tau})$

end if

else

RJ-step (4): Delete changepoint

 Delete a changepoint τ' , propose new segment parameters $\tilde{\theta}$

 Compute acceptance ratio

$$\alpha(\tilde{s}, s) = \min \left\{ 1, \frac{p(\mathbf{y}|\tau', \tilde{\theta}, \tilde{M}) p(\tau') p(\tilde{\theta}) p(\tilde{M}) q(\tau|\tau') q(\theta|\tilde{\theta}) q(M|\tilde{M})}{p(\mathbf{y}|\tau, \theta, M) p(\tau) p(\theta) p(M) q(\tau'|\tau) q(\tilde{\theta}|\theta) q(\tilde{M}|M)} |J| \right\}$$

 where $|J| = \frac{1}{2}$ and $q(\tilde{M}|M)$ is symmetrical to $q(M|\tilde{M})$.

 Generate $v \sim \text{Uniform}[0, 1]$

if $v \geq \alpha$ **then**

 Accept new state

$s = (M, \tilde{\theta}, \tau')$

end if

end if

end for

Remove first B samples and keep every T^{th} sample

Return: $\{s\}_{t=1}^T$

2.4 Locating change in airway dilatation

For airways affected by IPF, dilatation starts at the distal point and progresses in the proximal direction [7]. Therefore, we topologically can assume that each affected airway undergoes a single changepoint from which dilatation starts. To locate such unique changepoint, we consider the posterior probability of the changepoint $p(\tau|\mathbf{y})$, and perform the following post-processing steps.

On each airway track, the proximal region is surrounded by cartilage [13]. As the airway track loses cartilage support, the geometry of the airway changes and results in a changepoint. Since such biological changepoint is independent of the disease state and occurs prior to dilatation, we eliminate it by discounting the most proximal peak in the posterior distribution $p(\tau|\mathbf{y})$. We then selected highest peak on the modified posterior $p(\tau|\mathbf{y})$ as the final estimation for the starting point of dilatation.

3 Evaluation & Results

We evaluated our proposed method with two experiments. Firstly, we used a set of healthy airways, in which we had simulated dilatation to quantify the accuracy of our method and compared the results with conventional changepoint methodologies. Secondly, we assessed the clinical utility of our method on a dataset of airways affected by IPF. We compared our labels with those from two experienced thoracic radiologists. The properties of the images used in both experiments are displayed on Tab. 1.

Experiment	Patients	BL Voxel Size	FU Voxel Size	Airways	Time between scans
1	1	0.67, 0.67, 1.00	0.56, 0.56, 1.00	6	9M 6D
1	2	0.63, 0.63, 1.00	0.78, 0.78, 1.00	7	35M 6D
1	3	0.72, 0.72, 1.00	0.72, 0.72, 1.00	1	5M 22D
2	1	0.72, 0.72, 1.00	0.64, 0.64, 1.00	3	12M 5D
2	2	0.67, 0.67, 1.00	0.87, 0.87, 1.00	1	10M 24D

Table 1. Table of the image properties of voxel size, number of airways used and the time between scans. The airways were selected by a trained radiologist R1. All patients in Experiment 1 do not have IPF. Abbreviation: BL - Baseline Scan, FU - follow-up Scan, M - Months, D - Days.

3.1 Disease simulation

To quantitatively assess accuracy, a ground truth is required. To this end, we applied our point detection algorithm on augmented healthy airway series to simulate the airway dilatation caused by IPF. After obtaining written informed consent from 3 patients, a trained radiologist (R1) selected 14 pairs of healthy airways in both baseline and follow-up scans. The image properties are displayed on Tab. 1. They were acquired from different scanners and used different reconstruction kernels. The airways were pre-processed as described in Sec. 2.1 to produce a function of area change along the length of the airway. We interpret this output signal as the overall noise caused from normal biological changes and acquisition.

We modelled the change in dilatation with a logistic function $l = \frac{M}{1+e^{-k(x-\alpha)}}$. We interpreted M as the magnitude of dilatation and α as the point of dilatation. The parameters α are set such that the dilatation starts 10-40mm from the distal point in increments of 5mm. In addition, we set M to range from 0.3-3 in increments of 0.25. Finally we set $k = 0.5\text{mm}^{-1}$, in order to create an abnormal increase in cross sectional area. To simulate the dilatation on the airway; the logistic function was added to the area change of the healthy airway, as shown on Fig. 2. We applied every permutation of M and α on each of the 14 healthy airways. After the signal augmentation, we applied our proposed Bayesian changepoint detection algorithm.




Fig. 2. (A). An example of the area change of an airway remaining healthy on both baseline and follow-up scan. (B). A logistic function we constructed to simulate a dilatation due to IPF. (C). The augmented area change. The red line corresponds to our ground truth as the starting point of dilatation α . (D). The corresponding reconstructed airways. The images were acquired using Quan et al [11].

The proposed method was compared against two conventional methods. First, we use a basic thresholding method. For a given signal, we applied a 5mm moving mean average and thresholded the point at which the signal reached above the upper quartile from the right hand side. Secondly, we implemented the method based on Lavielle [8], in summary we consider K changepoint and these changepoints y_i , minimize the function: $J(K) = \sum_{r=0}^{K-1} \sum_{i=k_r}^{k_{r+1}-1} \Delta(y_i, y_{k_r:k_{r+1}-1}) + \beta K$, where β is modified such that the function finds less than K changepoints. For our paper, different $\Delta(y_i, y_{k_r:k_{r+1}-1})$ were evaluated and we found $\Delta(y_i, y_{k_r:k_{r+1}-1}) = y_i - \text{mean}(y_{k_r:k_{r+1}-1})$ gives the most accurate results. To replicate the post processing of our proposed method, we consider $K = 2$ possible changepoints. This takes into account the changepoint caused by the support cartilage. Finally, we set a minimum distance of 20mm. Once we acquired the changepoints, the most peripheral point was chosen as the point of dilatation. The implementation was performed through Matlab inbuilt function; `findchangepts`¹.

Fig. 3 shows the accuracy for each individual method as a heatmap. The metric used to quantify accuracy was the displacement between the ground truth and the changepoint given by the proposed method. A positive displacement (mm) corresponds to an overestimation of the ground truth towards the distal point. Each entry on the heatmap corresponds to the median average of all displacements over 14 airway pairs.

¹ <https://www.mathworks.com/help/signal/ref/findchangepts.html> last accessed on December 14, 2018.




Fig. 3. Heatmap showing the accuracy of each method. (Left). Thresholding method. (Middle). Method from Lavielle [8]. (Right). Our proposed method. The colour scale is the same on all of the heatmaps.

When the magnitude of dilatation is larger than $M > 0.75$, our proposed method achieves consistently higher accuracy than Lavielle [8]. We note that the accuracy gain in the peripheral regions of the airways at $\alpha = 10\text{-}30\text{mm}$ from the distal point are the most clinically relevant in IPF as lung parenchymal damage begins in the lung periphery and progresses proximally [7]. Furthermore on the same peripheral regions $\alpha = 10\text{-}30\text{mm}$, the baseline method showed systematic bias in accuracy towards the central airways. This was due to the baseline method being influenced by outliers from the longer expanses of normal airway regions. The proposed method uses the t-distribution as the likelihood thus making it robust to possible outliers within the data [10]. On the other hand, the proposed method suffers from poor accuracy below magnitudes of dilatation $M = 0.75$. However, in physical terms a dilatation of $M = 0.75$ corresponds to a percentage increase in cross sectional area of $e^{0.75} - 1 \approx 112\%$. This is within the range of normal biological change of the airways [2].

3.2 Application to Airways affected by IPF

We applied our method to airways affected by IPF. The purpose was to compare our measurement to the labels provided by a radiologist and compute the volume change of the identified diseased airway regions. For our dataset, we acquired 4 airway pairs from 2 patients after obtained a waiver for consent from the local Research Ethics Committee. All airways were judged by the radiologist R1 to be dilated as a consequence of IPF on baseline and to have visually worsened on follow-up imaging. Image properties are displayed in Tab. 1.

We compared the performance of our method against two trained thoracic subspecialist radiologists. Two radiologists R1, R2 identified the point of at which a given airway was seen to demonstrate increased dilatation on the follow-up CT scan. To assessed the reproducibility of manual labeling, each radiologist labelled the same airway twice through two different protocols. In the first method, the radiologists interrogated axial CT images. Using 2 separate workstations and the airway centreline, the radiologists identified the point on the centreline (on the follow-up scan) where the airway demonstrated definitive worsened dilatation.

For the second method, the radiologist compared the aligned reconstructed cross-sectional planes on baseline and follow-up scans. The radiologist then selected the slice where the airway had worsened when evaluated against the baseline scan. An example of both protocols are displayed on Fig. 4.




Fig. 4. Example views of images in two measurement protocols employed by radiologists to locate starting points of dilatation. (Left). Protocol 1 based on the axial slice. (Right). Protocol 2 based on the reconstructed cross sectional planes.

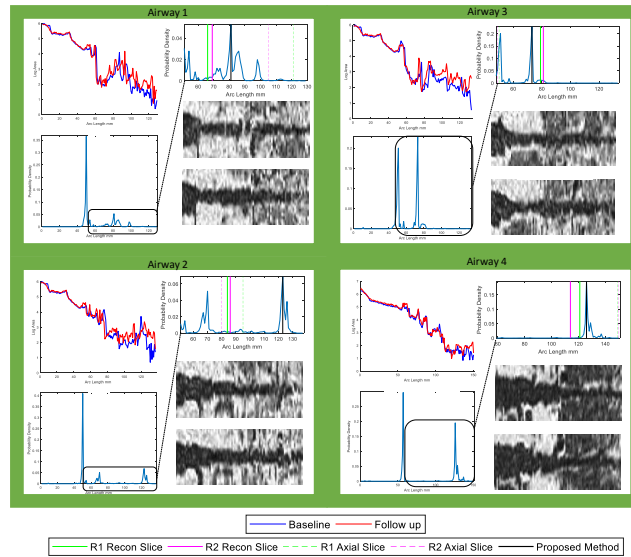


Fig. 5. The log cross sectional area and posterior distribution $p(\tau|\mathbf{y})$ for each of the four airway pairs. Airways 1,2,3 come from the same patient. In the magnified region (black) we compared the labels from our proposed method with the radiologist. Furthermore, we displayed the reconstructed slices within the magnified region.

Fig. 5 compares the predictions of our method with the labels from radiologists obtained in two different protocols. The results indicate that the predictions for Airway 1, 3 and 4 are within the range of the radiologists' labels. In the case of Airway 2, although our method based on the maximum peak overestimates with respect to the radiologists' predictions, the posterior distribution contains another equally probable peak that underestimates the radiologists' labels (see the second highest peak at 70mm), potentially indicating a more proximal point of dilatation. To test this, we delineated the boundary of the lumen on the reconstructed cross sectional slices at baseline in the neighbourhood of this peak,

and Fig. 6 shows the initial few slices (62-64mm). We observed that, when the delineated boundary from baseline was superimposed on the follow-up scan, the boundary is contained inside of the follow-up lumen with several lumen pixels are consistently outside from the boundary. Thus, this result indicates that the starting point of dilatation is more proximal than the labels from the radiologists.




Fig. 6. A row of three consecutive reconstructed slices in Airway 2 located on the arc length of 62-64mm. Each slice shows the airway lumen at baseline (left) and follow-up (right). The boundary delineation (red) from the baseline are superimposed on the follow-up scan. The blue arrows indicate pixels from the lumen outside the boundary.

To demonstrate the clinical utility of locating the starting point of airway dilatation, we compared longitudinal airway volume changes in diseased and healthy regions of each of the airway tracks. To find the volume of the airway, we considered the aligned signals, f_F, f_B . These signals are measurements of area against the airway arc length. Thus volume can be computed via the area under the curve. For our work, we used the trapezium rule to find the volume. Three volumetric regions were considered: (i) the entire airway $V_{c \rightarrow d}$ i.e. from the carina, c to the distal point, d , (ii) carina to the dilatation point t , denoted as $V_{c \rightarrow t}$. (iii) the dilatation point to the distal point, denoted by $V_{t \rightarrow d}$. Note that $V_{c \rightarrow d} = V_{c \rightarrow t} \cup V_{t \rightarrow d}$ and $V_{c \rightarrow t}$ does not overlap with $V_{t \rightarrow d}$. Tab. 2, shows the results of the percentage volume change. The volume change in $V_{t \rightarrow d}$ had greater sensitivity for selecting progressive airway dilatation in IPF than the volume change in the entire airway $V_{c \rightarrow d}$.

Airway	PVC of $V_{c \rightarrow d}$	PVC of $V_{t \rightarrow d}$	PVC of $V_{c \rightarrow t}$
1	2.6 %	32.9%	1.1%
2	3.2 %	129.7%	2.2%
3	2.6 %	47.4 %	-0.3%
4	7.4 %	48.4 %	7.1%

Table 2. The percentage volume change (PVC %) for each region of the airway.

4 Discussion & Conclusion

In this paper, we modelled changes in area along airway tracks as a time series with the purpose of detecting abnormal dilatation caused by IPF using Bayesian changepoint detection. Experiments on simulated data show that our model is able to detect the starting location of airway dilatation with superior accuracy than the relevant baseline methods. The results on the IPF longitudinal dataset display reasonable agreement with radiologists, while in one case indicating a more plausible location of dilatation, potentially missed by the experts.

Identifying changepoints and thereby calculating a change in airway dilatation over time could become a sensitive measure of IPF aggravation. This would form an important secondary endpoint in drug trials, where our measurements could indicate whether a new drug ameliorates disease progression better than existing medications. In the future, we hope to evaluate such utility of our proposed method on a larger cohort of IPF subjects. Furthermore, our method is applicable to other lung airway diseases, characterised by dilatation (such as cystic fibrosis) or other geometrical deformations, and such extensions remain valuable future work.

References

1. Siddhartha Chib and Edward Greenberg. Understanding the Metropolis-Hastings Algorithm. *The American Statistician*, 49(4):327–335, 1995.
2. Lee Gazourian, Samuel Ash, Emily E.K. Meserve, Alejandro Diaz, Raul San Jose Estepar, Souheil Y. El-Chemaly, Ivan O. Rosas, Miguel Divo, Anne L. Fuhlbrigge, Phillip C. Camp, Vincent T. Ho, Ami S. Bhatt, Hilary J. Goldberg, Lynette M. Sholl, and George R. Washko. Quantitative computed tomography assessment of bronchiolitis obliterans syndrome after lung transplantation. *Clinical Transplantation*, 31(5):e12943, 2017.
3. Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis, Third Edition*, volume 1. CRC Press, 2014.
4. Peter J Green and David I Hastie. Reversible jump MCMC. *Genetics*, 155(3):1391–1403, 2009.
5. Suicheng Gu, Carl Fuhrman, Xin Meng, Jill M. Siegfried, David Gur, Joseph K. Leader, Frank C. Sciruba, and Jiantao Pu. Computerized identification of airway wall in CT examinations using a 3D active surface evolution approach. *Medical Image Analysis*, 17(3):283–296, 2013.
6. J Jacob, B J Bartholmai, S Rajagopalan, M Kokosi, A Nair, R Karwoski, S L F Walsh, A U Wells, and D M Hansell. Mortality prediction in idiopathic pulmonary fibrosis: evaluation of computer-based CT analysis with conventional severity measures. *European Respiratory Journal*, 49(1):1601011, 2017.
7. Joseph Jacob and David M. Hansell. HRCT of fibrosing lung disease. *Respirology*, 20(6):859–872, 2015.
8. Marc Lavielle. Using penalized contrasts for the change-point problem. *Signal Processing*, 85(8):1501–1510, 2005.
9. T. Mikosch and C. Starica. Changes of structure in financial time series and the GARCH model. *REVSTAT Statistical Journal*, 2(1):41–73, 2004.
10. S J D Prince. *Computer Vision: Models, Learning, and Inference*. Cambridge University Press, 2012.
11. Kin Quan, Rebecca J. Shipley, Ryutaro Tanno, Graeme McPhillips, Vasileios Vavourakis, David Edwards, Joseph Jacob, John R. Hurst, and David J. Hawkes. Tapering analysis of airways with bronchiectasis. *Proceedings of SPIE*, 10574:105742G, 2018.
12. David Siegmund, Benjamin Yakir, and Nancy R Zhang. Detecting simultaneous variant intervals in aligned sequences. *The Annals of Applied Statistics*, 5(2A):645–668, 2011.
13. E R Weibel. *Morphometry of the Human Lung*. Springer, 1963.