

PAPER • OPEN ACCESS

Towards automated *in vivo* tracheal mucociliary transport measurement: Detecting and tracking particle movement in synchrotron phase-contrast x-ray images

To cite this article: Mark Gardner *et al* 2020 *Phys. Med. Biol.* **65** 145012

View the [article online](#) for updates and enhancements.



MR Safe
4D Phantom
for MRgRT

QUASAR MRI^{4D}
BY MODUS CA

LEARN MORE ►



PAPER

OPEN ACCESS

RECEIVED
25 November 2019REVISED
4 February 2020ACCEPTED FOR PUBLICATION
11 February 2020PUBLISHED
16 July 2020

Original Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Towards automated *in vivo* tracheal mucociliary transport measurement: Detecting and tracking particle movement in synchrotron phase-contrast x-ray images

Mark Gardner^{1,2} , David Parsons^{1,2} , Kaye Morgan^{3,4,5,6} , Alexandra McCarron^{1,2} , Patricia Cmielewski^{1,2} , Regine Gradl^{4,5,6} and Martin Donnelley^{1,2} ¹ Robinson Research Institute and Adelaide Medical School, University of Adelaide, Adelaide, Australia² Respiratory and Sleep Medicine, Women's and Children's Hospital, 72 King William Road, SA 5006 North Adelaide, Australia³ School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia⁴ Chair of Biomedical Physics, Department of Physics, Technical University of Munich, James-Frank-Str. 1, 85748 Garching, Germany⁵ Munich School of BioEngineering, Technical University of Munich, Boltzmannstr. 11, 85748 Garching, Germany⁶ Institute for Advanced Study, Technical University of Munich, Lichtenbergstr. 2 a, 85748, Garching, GermanyE-mail: mark.gardner@adelaide.edu.au**Keywords:** cystic fibrosis, mucociliary transport, particle tracking, cascade classifier, automation, synchrotron, Phase-Contrast x-raySupplementary material for this article is available [online](#)

Abstract

Accurate *in vivo* quantification of airway mucociliary transport (MCT) in animal models is important for understanding diseases such as cystic fibrosis, as well as for developing therapies. A non-invasive method of measuring MCT behaviour, based on tracking the position of micron sized particles using synchrotron x-ray imaging, has previously been described. In previous studies, the location (and path) of each particle was tracked manually, which is a time consuming and subjective process.

Here we describe particle tracking methods that were developed to reduce the need for manual particle tracking. The MCT marker particles were detected in the synchrotron x-ray images using cascade classifiers. The particle trajectories along the airway surface were generated by linking the detected locations between frames using a modified particle linking algorithm. The developed methods were compared with the manual tracking method on simulated x-ray images, as well as on *in vivo* images of rat airways acquired at the SPring-8 Synchrotron.

The results for the simulated and *in vivo* images showed that the semi-automatic algorithm reduced the time required for particle tracking when compared with the manual tracking method, and was able to detect MCT marker particle locations and measure particle speeds more accurately than the manual tracking method. Future work will examine the modification of methods to improve particle detection and particle linking algorithms to allow for more accurate fully-automatic particle tracking.

1. Background

Cystic fibrosis (CF) is a genetic condition that affects more than 30 000 people in the USA (Cystic Fibrosis 2019). It is caused by mutations in the CF transmembrane conductance regulator (*CFTR*) gene, which leads to dehydration of the Airway Surface Liquid (ASL) and a reduction in the efficiency of the Mucociliary Transport (MCT) system. Together these facilitate the production of thick sticky mucus that enables inhaled pathogens and foreign particles to accumulate in the lung. This results in a persistent cycle of infection and inflammation that leads to destruction of the lung tissue and eventually lung failure (Boucher 2019).

Altered ASL height and MCT behaviour are the first local airway surface changes that are responsible for the decline in lung health, however both can be challenging to accurately measure *in vivo*. Most pre-clinical methods for assessing MCT behaviour in live animals quantify mucus movement by depositing MCT

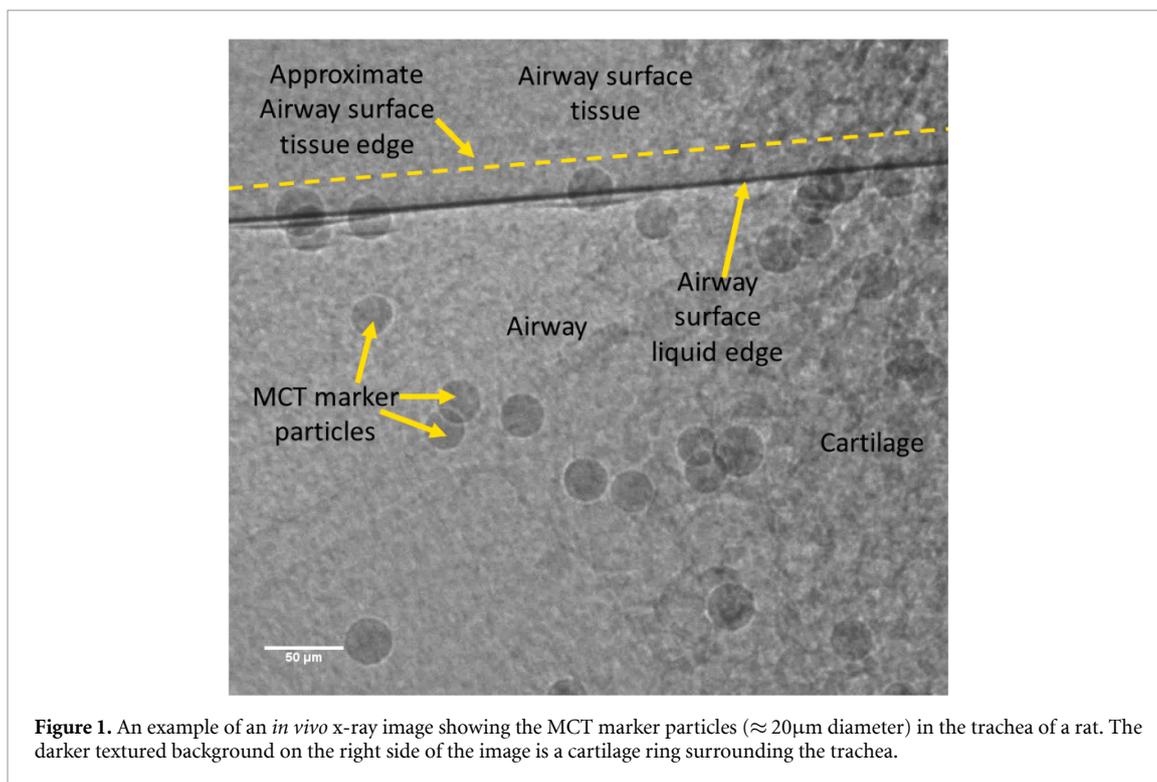


Figure 1. An example of an *in vivo* x-ray image showing the MCT marker particles ($\approx 20\mu\text{m}$ diameter) in the trachea of a rat. The darker textured background on the right side of the image is a cartilage ring surrounding the trachea.

marker particles onto the airway surface, and tracking their motion over time. In our implementation, micron-sized tracking particles are deposited onto the mucus, and their location is measured at high magnification using synchrotron phase-contrast x-ray imaging (S-PCXI) (Endrizzi 2018). S-PCXI captures how the x-ray light that passes through each particle is re-directed away from the centre of each particle, resulting in an image where each particle appears dark with a bright outline. An example image showing $20\mu\text{m}$ alumina particles in the trachea of a live rat is shown in figure 1. This method has been tested *in vivo* in mice (Donnelley *et al* 2014a, Donnelley *et al* 2014b, Donnelley *et al* 2016) and pigs (Donnelley *et al* 2019), and *ex vivo* in pigs (Donnelley *et al* 2017), sheep (Donnelley *et al* 2017) and rats (Gardner *et al* 2019). A review describing the experimental methods in greater detail has been published (Morgan *et al* 2019). Unlike other marker-transport-based methods that quantify the behaviour and speed of large groups of particles (Rogers *et al* 2018, Hua *et al* 2010), the S-PCXI method allows the velocity of each individual particle to be measured (Donnelley *et al* 2017, Donnelley *et al* 2019), enabling local MCT variability to be assessed.

Prior to the development of the algorithms described in this manuscript the particle motion in the S-PCXI datasets was manually tracked using custom Matlab (R2019a, Mathworks, USA) software that presents an observer with a short sequence of image frames in which particles should be tracked. The observer manually clicks to identify the location of the same particle in each of the frames, and repeats the process until every moving particle in that sequence is identified and tracked. This is repeated for every time-point and every subject in the data set. Depending on the number of moving particles, this manual tracking process can take more than a few days to complete, is somewhat subjective, and can be inaccurate especially for slow moving particles.

We propose developing methods to automate sections of the particle identification and tracking would make the particle tracking process faster and more accurate. However, the poor signal to noise ratio (SNR) of the captured S-PCXI images, combined with overlying features such as the tracheal cartilage (see figure 1), makes automatic particle tracking difficult. There have been previous attempts to detect MCT particles using traditional pattern detection algorithms (Jung *et al* 2017) and convolutional neural networks (Jung *et al* 2019), as well as to track their motion (Jung *et al* 2019). However, these previous attempts were designed specifically for *in vivo* images acquired from mouse nasal airways instead of the rat trachea, and were performed using particles with a different composition to the particles used in more recent methods (Gardner *et al* 2019, Donnelley *et al* 2019), and thus appear different under S-PCXI.

Methods have been developed for accurate and reliable automatic particle tracking for the purposes of single particle tracking to study sub-cellular dynamics (Jaqaman *et al* 2008, Newby *et al* 2018, Sbalzarini and Koumoutsakos 2005, Chenouard *et al* 2014). Additionally, Newby *et al* demonstrated a case where an automatic particle tracking algorithm was faster than manual tracking (Newby *et al* 2018). However, the

particle detection problem for MCT measurement using S-PCXI differs from more common particle detection problems in two areas; particle identification and particle tracking.

Firstly, in most other tracking applications the image contrast is high, with the particles that are being tracked providing a significantly different intensity to the background intensity, making their identification relatively easy. This is different to the example situation presented in figure 1, which shows that the pixel values within the MCT marker particles are similar to the pixel values outside the particle, and the particles are primarily identified by their edges. This is because S-PCXI enhances the boundaries between materials with different x-ray refractive indices, even when they are weakly absorbing due to their extremely small mass. This means that the particle profiles are substantially more complex than in situations where high pixel values correspond to the particles and low pixel values belong to the background. As a result, non-trivial object detection methods are needed to identify the MCT marker particles.

Secondly, previous examples of automated particle tracking were able to utilise a high frame rate (Jaqaman *et al* 2008, Newby *et al* 2018), which reduces the potential distance that each particle can move between frames. This makes the process of linking the locations of the identified particles into trajectories relatively easy. However, previous examples of the S-PCXI MCT measurement method describe how one image is acquired per breath because the available x-ray flux places a lower bound on the exposure time required for maintaining adequate image SNR. For anaesthetised mice and rats—which breathe at approximately 80 br/min and 120 br/min respectively—this means that the time between sequential images can be 0.5 – 0.75 sec, assuming one image is acquired per breath. This timing creates challenges for automatic tracking of the MCT marker particles, because particles with similar appearance can potentially move a large distance between frames, and hence algorithms to link particle locations are likely to generate many false particle linkages. Accurate particle tracking is further exacerbated by substantial background motion that naturally results when imaging the respiratory system of a live anaesthetised animal.

In this paper we describe a novel application of image processing methods that was developed to aid in the tracking of the MCT marker particles in S-PCXI images. Two different methods are presented; a fully-automatic particle tracking method and a semi-automatic method, in which the user can choose to use the automatic particle tracking results or manually track the particles. The automatic particle tracking process is split into two sections: particle identification and particle linking between consecutive frames. The automatic and semi-automatic tracking methods are compared with manual tracking results on a set of artificial images designed to simulate simple particle movement, and are also applied to real-world S-PCXI images taken from live anaesthetised rats.

2. Methods

2.1. Imaging

2.1.1. *In vivo* imaging

Imaging was conducted at the BL20XU beamline of the SPring-8 synchrotron radiation facility in Japan. Normal Wistar rats were anesthetized and the MCT marker particles (diameter 22 μm , composed of Al_2O_3 with a COOH coating) were dispersed into the airway using a miniature bronchoscope inserted briefly into the trachea via the mouth (McIntyre *et al* 2018). The anaesthetised animal was then placed into the x-ray imaging enclosure, and allowed to breathe naturally (no artificial ventilation was applied). To reduce motion artefacts caused by respiration, images were captured between breaths at end-expiration, making the temporal resolution of the images the same as the breathing rate of the animals. More detailed information on the experimental setup is provided in (Morgan *et al* 2019). All animal studies were performed in accordance with protocols approved by the University of Adelaide (M-2017-113) and SPring-8 Synchrotron animal ethics committees.

2.1.2. Simulated MCT S-PCXI image sequences

Prior to testing the particle detection and tracking methods on the *in vivo* images, they were tested on artificial image sequences that simulate particle movement. These images were generated to enable accurate comparisons of the manual, automatic and semi-automatic particle tracking methods, as the ground-truth locations and velocities of the particles were known. In contrast, there is no ground-truth for the real *in vivo* image sequences because the particle locations must be manually identified, making the accuracy of the manual and automatic methods difficult to compare.

The creation of each image was divided into two phases; creating the image background, and inserting the particles into the image. In total, 10 sets of images were created, with each set containing 50 images (500 images in total). The particle backgrounds for each image set were generated using randomly selected *in vivo* S-PCXI images from which the user selects the largest possible area within the image that contained no particles. This area was then copied and flipped horizontally and vertically and used as tiles to create an

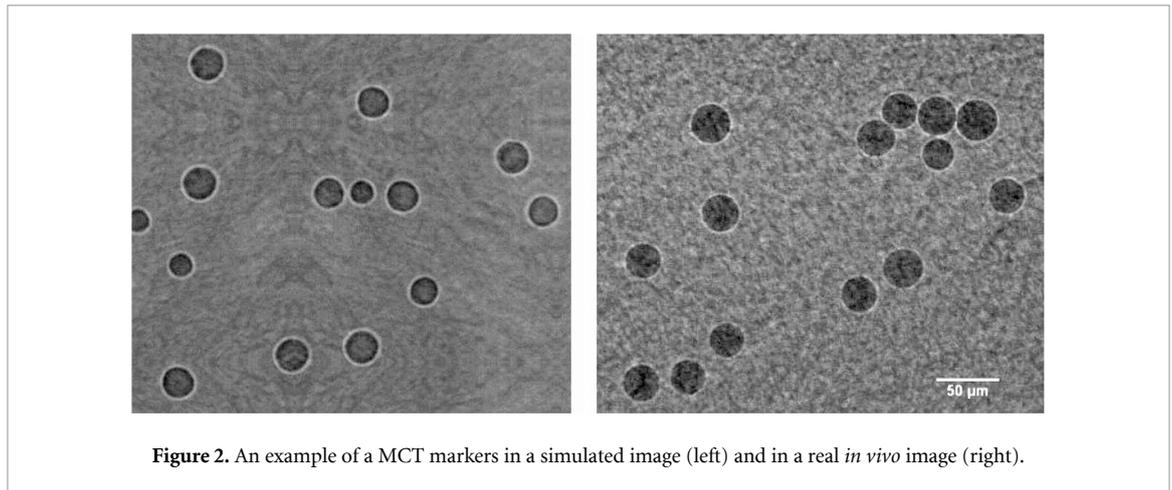


Figure 2. An example of a MCT markers in a simulated image (left) and in a real *in vivo* image (right).

image with the same pixel dimensions as the original *in vivo* S-PCXI image. These sections of the *in vivo* images were used so that the performance of the particle detection algorithms on the simulated data would better represent the performance when applied to the *in vivo* data.

To simulate the particle locations, images of particles by themselves (randomly selected from a pool of five different particles) were captured by depositing the particles onto x-ray translucent polyimide tape. Individual particles were then cropped from these images and were added to the background images generated using the formula:

$$im_{new} = im_{background} + [im_{particle} - median(im_{background}) - median(im_{particle})]. \quad (1)$$

The number of particles added to the images was chosen to mimic the *in vivo* data. An example of a static artificial image is shown in figure 2, and appears visually consistent with a real *in vivo* image.

To simulate particle movement by MCT through a sequence of images, some of the particle locations were changed between consecutive image frames. Not all particles move in a typical *in vivo* animal experiment, so to simulate partial particle movement, a number was randomly generated (in the range [0,1]) for every particle in every frame. Particles were set to begin moving within that frame if that number was greater than a defined threshold (set as 0.8 for these results). The speed at which the particles move in the artificial image sequences was also randomly selected, with a maximum possible speed of 3 mm/min based on preliminary results from manual tracking of the *in vivo* data. The direction of movement was also randomly assigned for each image set, with all particles in a given image set programmed to move in the same direction. If the number of particles was lower than a defined threshold (set as 30) then new particles were added to the image sequence.

2.2. Particle detection

The effectiveness of the particle detection algorithms was tested on 20 randomly selected frames. For the simulated data the particle detection algorithms were compared to the ground truth, and for the *in vivo* data the particle locations from the automatic particle detection algorithms were compared with the manually detected particle locations. For the different particle detection algorithms, the sensitivity, false positive rate and false negative rate were calculated, as well as the average time taken to analyse a frame. The particle locations were automatically detected by training cascade classifiers, and the results were compared to a circular Hough transform (CHT) method, a more traditional method for detecting circular objects.

2.2.1. Cascade classifiers

The MCT marker particles were detected in each image using a Cascade classifier (Viola and Jones 2004) implemented in Matlab. This classifier was chosen as it is able to quickly and reliably identify objects from an image. Two different Cascade classifiers were trained using different sets of features. These features were simple Haar type features (Viola and Jones 2004) and local binary patterns (LBP) (Ojala *et al* 2002). For both classifiers the number of stages, false positive rate, and true positive rate were set to 30, 0.05 and 0.995, respectively.

To train the classifier one frame was randomly selected from each of the *in vivo* specimens. From each of these frames, regions of interest (ROIs) containing particles were manually identified and used as examples of positive images. A set of negative images were obtained from image regions not containing any particles,

so that the negative images contained similar background features to improve the detection accuracy of the classifiers. A total of 5451 positive images and 10902 negative images from the 17 rats were used for each stage to train both Cascade classifiers.

To reduce the number of false positives detected, both the Haar and LBP classifiers were used to estimate particle locations in each frame, and any common points between both classifiers was taken as the final particle locations.

2.2.2. Comparison to circular Hough transform

The methods developed here were compared with a more traditional approach for detecting circular objects: A smoothing step followed by a CHT. The smoothing step involved applying two Gaussian smoothing functions (Nixon and Aguado 2012). These two functions were applied to each image, with σ values of $r\sqrt{2}$ and $r(\sqrt{2} - 1)$ where r is an average expected radius of the particles, which was set to 25 pixels. The difference between these images was then taken as a form of a band-pass filter to attempt to remove features that were not particles. From this image, a CHT was applied to identify the circular particles.

2.3. Particle linking

Once the location of all of the particles was determined in every image, the next step was linking the individual particles across the multiple images in the sequence, allowing the path of each particle to be measured. This section describes the previous manual linking method, as well as new automatic and semi-automatic particle linking methods.

2.3.1. Manual tracking method

As noted in section 1, the previous method for quantifying deposited particle MCT behaviour involved manual particle tracking (Donnelley et al 2014b). In that process, the user was shown ~ 15 seconds worth of consecutive frames from one specimen that were chosen from specific time-points throughout the imaging experiment. The user then manually identified the location of each moving particle over these selected frames by clicking on the centre of the particle. Each click to identify a particle advanced to the next frame in the sequence, allowing the particle trajectory to be followed. Once a single particle was tracked its location was marked on each frame of the sequence, and the process was repeated for the remaining particles. The user was instructed to only identify moving particles, so the position (and hence the velocity) of the stationary particles was not included in these MCT measurements. Once every moving particle was identified, either the next time-point or sample was displayed. To attempt to prevent any observer bias, the order of the specimens from which the frames were selected and the time-points at which the data was acquired were randomised, and the user was blinded to both. This meant that the user could not determine which sample (and therefore the treatment it may have received) or time-point (the time post-treatment delivery, and the magnitude of the treatment effect on MCT) a particular sequence of images belonged to. The distance the particle moved in between frames was converted from pixels to mm to calculate the distance travelled, then the elapsed time between the frames was used to calculate the MCT rate (in mm/min).

2.3.2. Automatic method

The particle locations for each frame were detected using the cascade classifiers described above. The next step was to link the particle positions in consecutive frames using a method based on work by Jaqaman et al (2008). Briefly, they describe that three possible scenarios for linking particles between frames must be taken into account: (1) finding a matching particle between consecutive frames, (2) the particle has disappeared between frames, or (3) a new particle has appeared. A cost matrix C was constructed showing the possible scores for linking the detected particles in frame t to the particles in frame $t + 1$. The structure of the matrix is shown below:

$$\begin{array}{c}
 \text{Frame}_t \backslash \text{Frame}_{t+1} \\
 \begin{array}{c}
 1 \\
 2 \\
 \vdots \\
 n_t \\
 - \\
 1 \\
 2 \\
 \vdots \\
 n_{t+1} - 1 \\
 n_{t+1}
 \end{array}
 \end{array}
 \left[\begin{array}{cccc|cccc}
 1 & 2 & 3 & \dots & n_{t+1} & 1 & 2 & \dots & n_t \\
 l_{11} & l_{12} & l_{13} & \dots & l_{1n_{t+1}} & d & X & \dots & X \\
 l_{21} & l_{22} & l_{23} & \dots & l_{2n_{t+1}} & X & d & \dots & X \\
 \vdots & \vdots & \vdots & & \vdots & \vdots & & \ddots & \\
 l_{n_t 1} & l_{n_t 2} & l_{n_t 3} & \dots & l_{n_t n_{t+1}} & X & X & \dots & d \\
 - & - & - & - & - & + & - & - & - \\
 b & g_1 & X & \dots & X & & & & \\
 X & b & g_2 & \dots & X & & & & \\
 \vdots & & \ddots & \ddots & & & & & \\
 X & X & \dots & b & g_{n_{t+1}-1} & & & & \\
 g_{n_{t+1}} & X & X & \dots & b & & & &
 \end{array} \right]$$

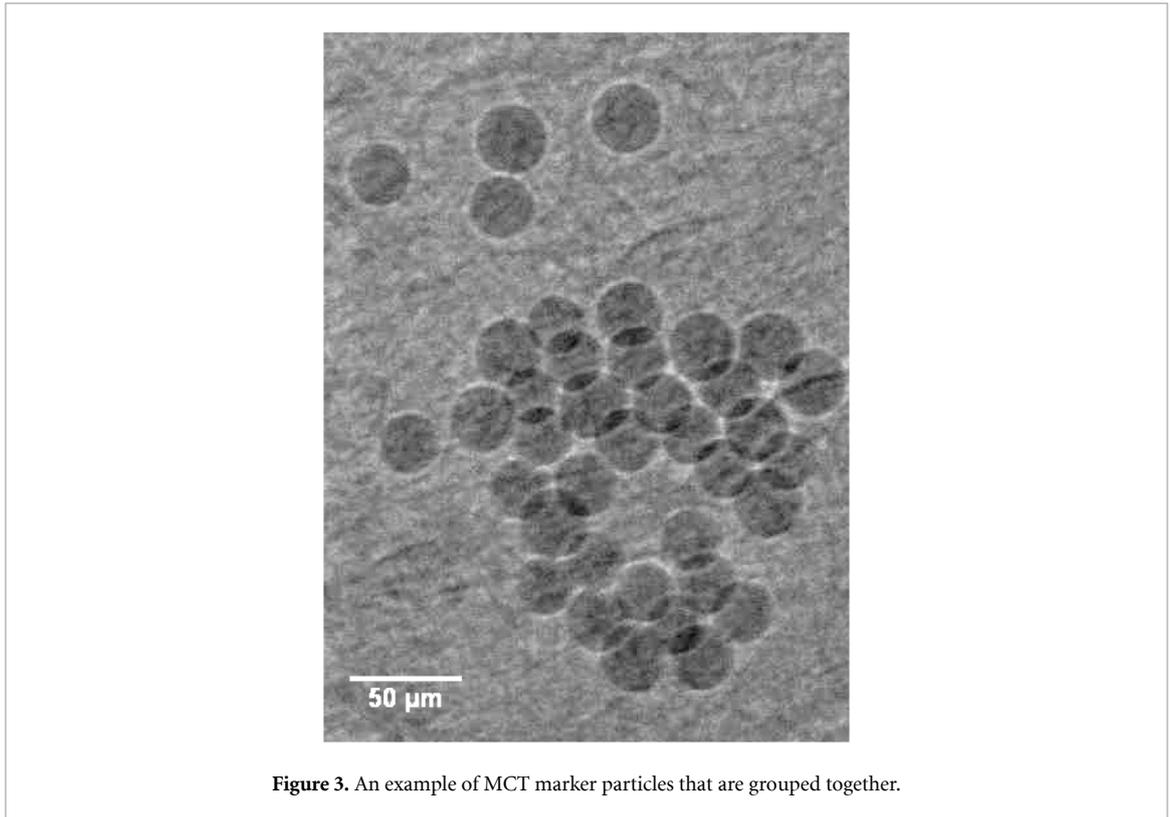


Figure 3. An example of MCT marker particles that are grouped together.

where l_{ij} is the distance between the adjusted location of particle i in frame t and particle j in frame $t + 1$, b and d are the scores for new particles appearing in frame $t + 1$ and disappearing in frame t respectively. Unlike the method used by Jaqaman *et al* (2008), an additional term g_j was added to the cost matrix C , which was defined as the minimum distance between particle j in frame $t + 1$ and all particles that were classified as disappearing in frame $t - 1$. This term was introduced as the particle detection accuracy shown in table 2 meant that sometimes particles were not detected in every frame, and some particles were incorrectly classified as ‘disappeared’ when they had in fact just not been detected. This additional term assisted in preventing particles being incorrectly linked as a results of a particle being misdeteched in a previous frame. This scenario is more common in this specific application as the MCT marker particles had a tendency to group together as shown in figure 3, greatly reducing the distance between particles.

Once the cost matrix C had been constructed the particle locations were linked by solving the linear assignment problem (LAP) for C with dimensions $P \times P$:

$$\hat{A}_{\text{argmin}} = \sum_{i=1}^P \sum_{j=1}^P A_{ij} C_{ij} \quad (2)$$

such that

$$\sum_{i=1}^P A_{ij} = 1, \sum_{j=1}^P A_{ij} = 1. \quad (3)$$

The bottom right section of matrix C was designed to allow for solving of the matrix using the LAP. The solution to the LAP problem was calculated using the Jonker–Volgenant algorithm (Jonker and Volgenant 1987). This algorithm was chosen as it is able to solve the LAP problem faster than traditional methods, including the Hungarian method (Munkres 1957), especially when a large number of particles are detected leading to a large cost matrix.

In this application a small component of the detected particle motion was caused by movement of the whole sample. Here, that effect results from factors such as respiration and small muscle movements from the anaesthetised animal, which are more pronounced when working at such high magnifications. To compensate for this movement error it was assumed that the change in pixel position of each particle between frames was represented by

$$\Delta_{\text{Measured}} = \Delta_{\text{Actual}} + \Delta_{\text{Frame}} \quad (4)$$

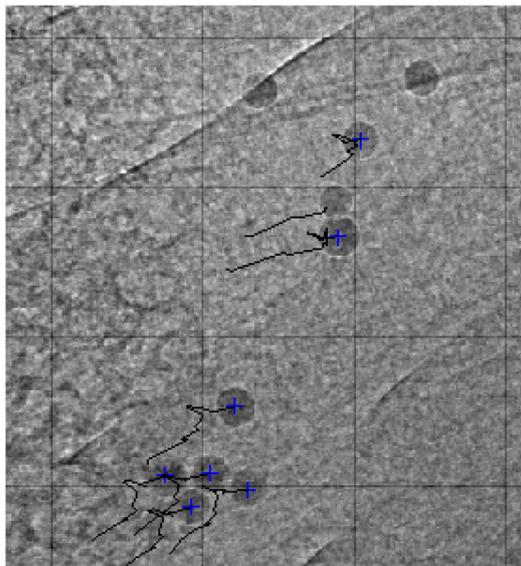


Figure 4. An example screen during semi-automatic particle tracking. The particle paths for any moving particles are shown as a black line, and the locations of those particles for that particular frame are shown as blue crosses. Note the paths of the two particles at the top of the image are not shown as they remained stationary for that time-point. Also note the one black line with no blue cross at the end indicating a particle that has been tracked in previous frames, but cannot be found in the current frame.

where $\Delta_{Measured}$ is the measured difference in particle location for each particle between successive frames, Δ_{Actual} is the actual movement of each particle between frames, and Δ_{Frame} is the component of the measured movement caused by the frame/sample moving. Whilst the values of $\Delta_{Measured}$ and Δ_{Actual} were unique to each particle, the value of Δ_{Frame} should be relatively constant for every particle in the frame because all particles were affected in a similar way. Hence the value for Δ_{Frame} for each frame was estimated by taking the median value of $\Delta_{Measured}$ for all particles in each frame, which allows Δ_{Actual} to be estimated.

2.3.3. Semi-automatic tracking method

As described in the imaging section, *in vivo* images were acquired between each breath, which for the rats that were being measured was approximately every 0.75-1 seconds. Given this relatively poor temporal resolution between images, the presence of artifacts and overlying features in the x-ray images (shown in figure 1), as well as the tendency for the particles to group together (shown in figure 3), the false detection of particles and the creation of false particle tracks could introduce significant errors into the results from the automatic particle tracking. Hence, a semi-automatic method was developed to help overcome these problems. This approach utilised the best of the automatic and manual tracking methods, by using the automatic tracking results to assist the manual tracking process.

In this semi-automatic method, the user was shown ~ 15 seconds worth of consecutive frames in the same manner as the manual tracking method, which were overlaid with the detected particle paths and locations from the automatic tracking method. An example of this is shown in figure 4. The user was then able to select particle tracks to save for some particles and manually enter particle positions for others, allowing the automatic method to track the trivial particle paths, and the manual method to track the more complex or occluded particle paths. Whilst the semi-automatic method still requires some level of user input, removing the need for the user to input the location of every moving particle for every frame can significantly reduce the time needed to input the particle location, in comparison to the manual tracking method used previously.

3. Results

3.1. Particle detection accuracy

3.1.1. Simulated data

The different particle detection methods described in section 2.1 were applied to all of the simulated images. The results from the different methods are shown in table 1, and demonstrate that there was no significant difference in the sensitivity or the false negative rates (FNR) of the different particle detection methods. However the false positive rates (FPR) of the particle detectors using the cascade classifiers was significantly lower than the particle detection method using the CHT. The FPR was lowest when the common particle locations of the LBP and Haar cascade classifiers were used together. Importantly, this combination did not

Table 1. Results of particle detection accuracy on simulated data. Results shown are mean (std dev). * $p < 0.05$, two-way ANOVA vs CHT with Tukey correction for multiple comparisons.

	Sensitivity Rate (%)	False Positive Rate (%)	False Negative Rate (%)	Detection Time (s)
CHT	98.9 (1.46)	41.55 (17.16)	1.10 (1.46)	2.56 (0.21)
LBP	97.15 (2.47)	30.88 (14.51)*	2.85 (2.47)	2.24 (0.28)
Haar	97.76 (2.15)	17.86 (10.23)*	2.24 (2.15)	4.63 (1.38)
LBP & Haar	97.20 (2.41)	5.38 (5.12)*	2.80 (2.41)	6.88 (1.23)

Table 2. A comparison of how effective each particle detection algorithms is at correctly identifying particles. Results shown are mean (std dev). * $p < 0.0001$, † $p < 0.01$, two-way ANOVA vs CHT with Tukey correction for multiple comparisons.

	Sensitivity (%)	False Positive Rate (%)	False Negative Rate (%)	Detection Time (s)
CHT	86.89 (17.45)	65.69 (21.29)	13.12 (17.45)	3.88 (0.33)
LBP	84.57 (5.70)	35.59 (18.96)*	23.93 (11.74)†	1.36 (0.34)
Haar	83.38 (9.61)	21.94 (14.49)*	16.62 (9.61)	3.73 (1.75)
LBP & Haar	79.71 (9.23)	5.75 (5.74)*	20.29 (9.26)	7.14 (3.16)

result in a significant reduction in the particle detection sensitivity. This shows that for the simulated data, using the common locations of the trained LBP and Haar cascade classifiers resulted in superior particle detection.

3.1.2. *In vivo* images

The results in table 2 show how effective each particle detection algorithm was at identifying particles in the *in vivo* images. Whilst the sensitivity of the CHT method was not significantly different to the other particle detection methods, it had a significantly higher FPR than the classifier-based particle detection methods ($p < 0.0001$). The range of values in the Sensitivity, FPR, and FNR was also larger than the other detection methods, indicating that the CHT method was less consistent between different frames than the other particle detection methods.

For both the LBP and Haar methods the FPR was high (35.59% for LBP and 21.94% for Haar), even though it was significantly lower than for the CHT method. By combining the information from the LBP and Haar methods (LBP & Haar method), the FPR was significantly reduced to 5.75% ($p < 0.0001$) without a significant reduction in the particle detection sensitivity ($p > 0.24$). However the process of running both detection methods meant the LBP & Haar method had the largest average processing time when compared with the other particle detection methods.

3.2. Particle linking results

3.2.1. Simulated data

3.2.1.1. Particle tracking error

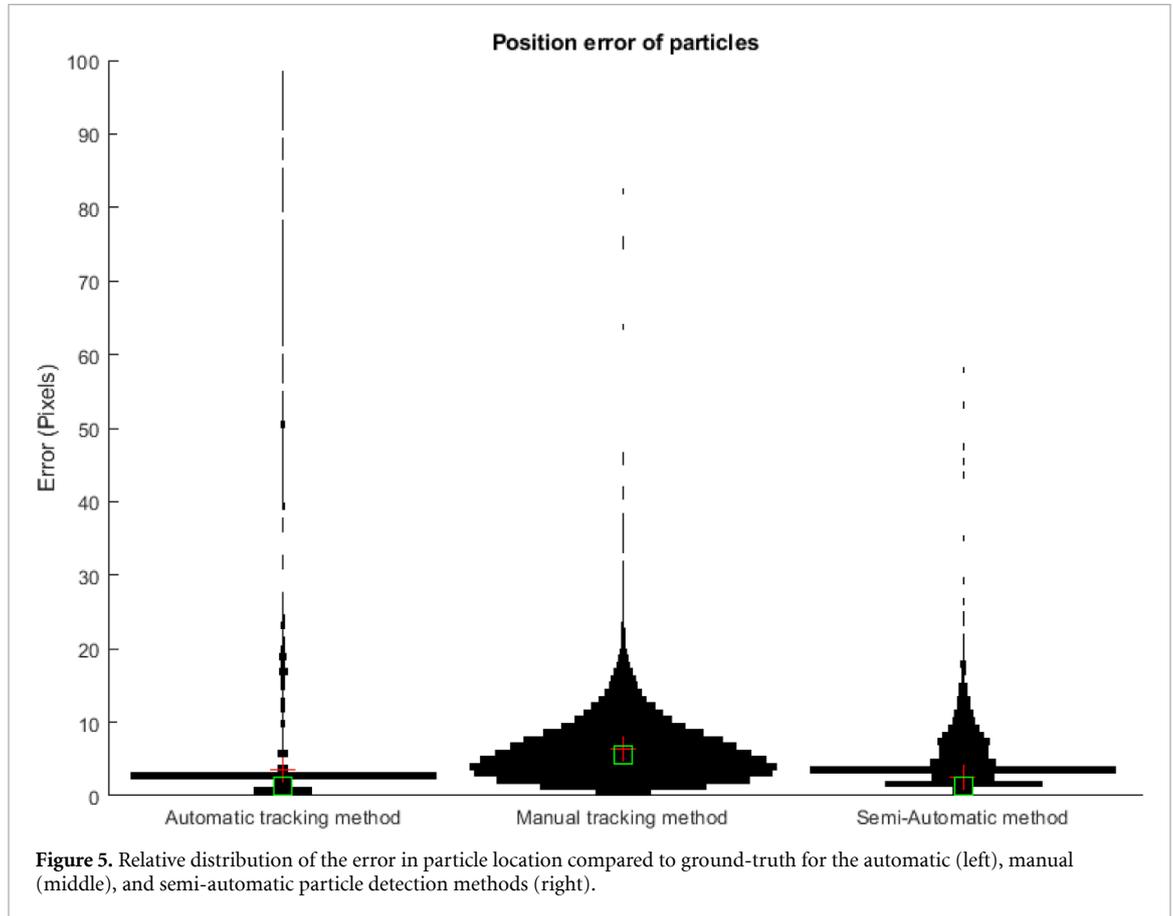
The results from the linking algorithm on the simulated particle movement are summarised in table 3. For the simulated data, the semi-automatic tracking reduced the time taken for the user to analyse the data by 34% when compared with the manual tracking, demonstrating how combining the automatic and manual particle tracking methods can reduce the time taken to analyse the data. Table 3 also shows that while the fully-automatic particle detection method requires no input from the user to track the particle locations, it was not able to track the particle locations as effectively as the manual or semi-automatic particle tracking methods.

For the simulated data, the automatic tracking algorithm was able to track 71% of the particle paths across all 10 frames for each time-point, and was on average able to track each particle across 83% of the frames for each time-point. In contrast, the manual and semi-automatic particle tracking methods were able to track 99% and 98% of particles across all 10 frames for each time-point, respectively. Hence, although the automatic particle detection method is able to track the particle locations without any input from the user, the accuracy of the particle tracks generated by the automatic particle detection method is significantly lower than the manual and semi-automatic particle detection methods.

The distribution of the accuracy of the particle detection and subsequent linking algorithm in detecting the particle location is shown in figure 5. This figure shows that the majority of the particle locations detected by the automatic method (left) are more accurate than the manual detection method (middle), which is also shown by the lower median particle location error in table 3. This is likely due to human error in locating the center of each particle. However, the automatic method is also more likely to have particle

Table 3. Summary of results from each method when applied to the simulated data.

	Manual	Automatic	Semi-automatic
Manual tracking time (hrs)	11.25	0	7.42
Mean particle location error (pixels)	9.70	6.90	3.01
Median particle location error (pixels)	5.48	1.41	1.41
Percentage of complete tracks (%)	99.31	70.92	98.67
Mean track completion rate (%)	99.70 (4.35)	83.27(30.46)	99.63 (4.24)



locations with significantly higher errors (i.e. larger than 40 pixels) than the manual method, which is shown by the distribution in figure 5 and the larger mean particle location error in table 3. These larger errors are caused by tracks being incorrectly assigned to particles that come close to each other.

The semi-automatic particle tracking method also was able to detect the particles more accurately than the manual tracking method, which is shown in 5 and by the lower median particle location error in 3. Unlike the automatic particle tracking method, the semi-automatic method had a reduced number of large particle location error values, which was a result of the user being able to correct any obviously false particle tracks that cause the large particle location errors for the automatic detection method. This results in the semi-automatic method having the lowest mean particle location error of the three methods.

3.2.1.2. Average particle speed

As mentioned previously, one desired outcome variable that has previously been extracted from the detected particle locations and tracks is the mean particle speed over all frames for each time-point (Donnelley *et al* 2014a, Donnelley *et al* 2017, Donnelley *et al* 2019, Gardner *et al* 2019). As such the mean particle speed was compared for the different particle detection methods and is shown in figure 6. As with the previous studies, only the speeds of particles classified as moving were measured. For the automatic methods a particle was classified as moving if the total displacement from the first to last frame was larger than 50 pixels, which was an empirically determined threshold based on observation of the simulated particle movement. For all detection methods a minimum particle speed threshold was used to only measure particle speeds when each particle was moving. The minimum threshold was set to 5 pixels per frame, which corresponds to approximately 0.2 mm/min, which was another empirically determined threshold.

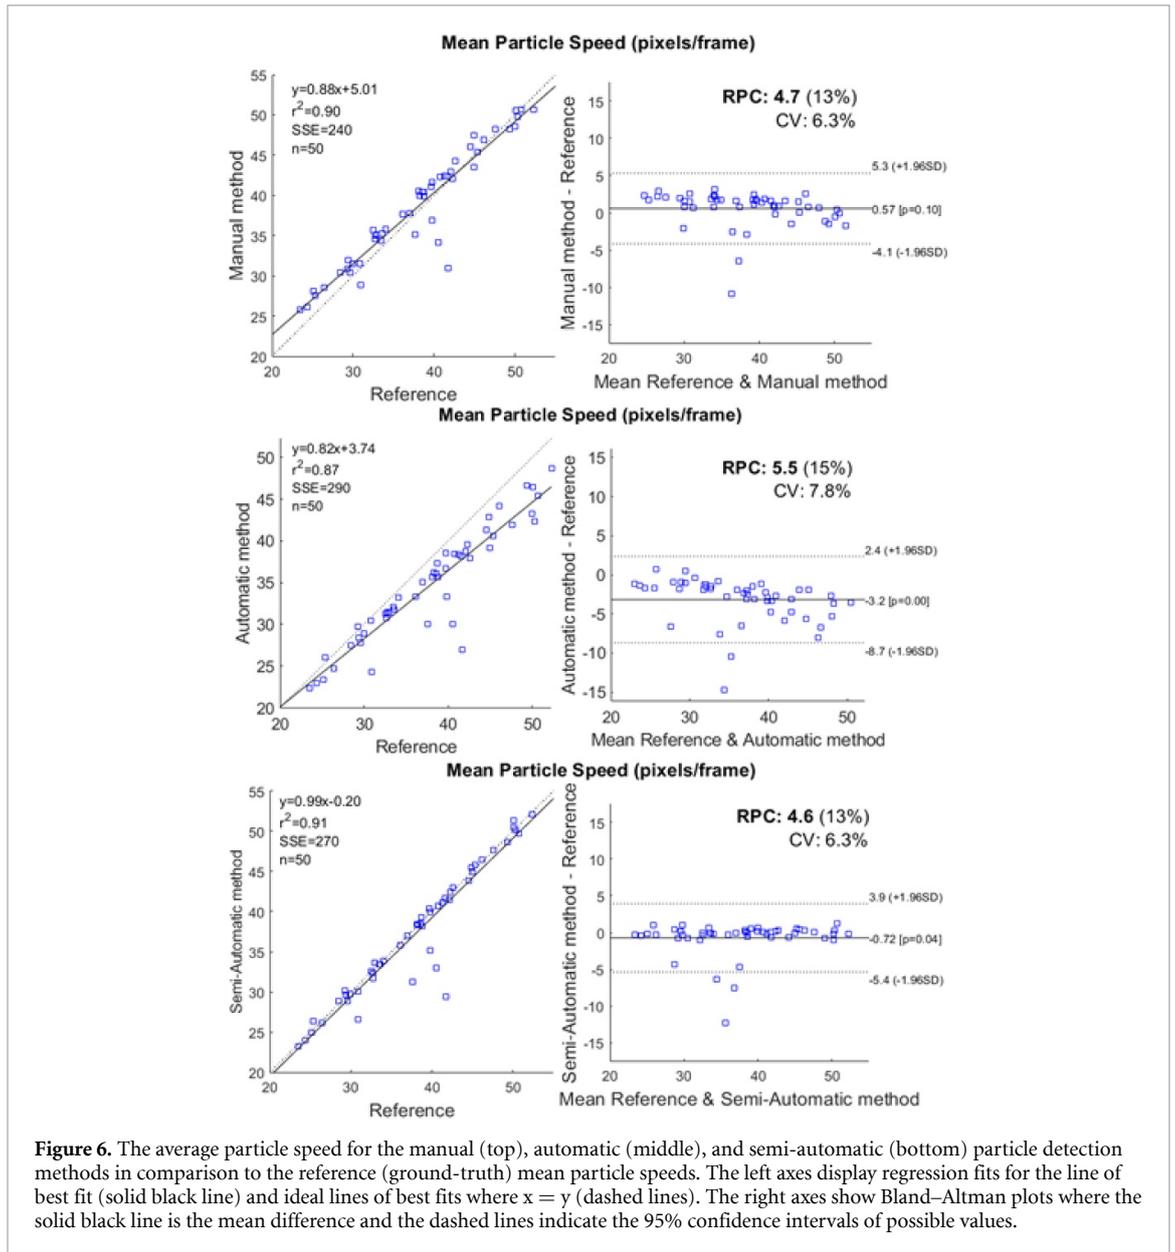


Figure 6 shows that the manual and semi-automatic methods are able to produce accurate mean particle speeds, for all possible mean particle speeds produced by the simulated particle movement. This is shown by the high r^2 values and the low repeatability coefficient (RPC) values. In contrast, the mean particle speeds from the automatic particle tracking method became less accurate for larger true mean particle speeds, which is shown by the diverging line of best fit (solid black line) and ideal line of best fit (dashed line).

3.2.2. *In vivo* data

3.2.2.1. Mean particle speed

For the *in vivo* images the mean particle speed was also calculated using the different particle tracking methods. As with the simulated data, only the moving particles were used for the mean particle speed calculations. For the automatic and semi-automatic method a particle was classified as moving if the total displacement over the sequence of frames was larger than 50 pixels (corresponding to a distance of $\approx 27 \mu\text{m}$). Minimum and maximum particle speed thresholds of 0.2 and 6 mm/min, respectively, were also applied for the *in vivo* MCT data, based on previously published methods (Donnelley *et al* 2017).

The mean *in vivo* particle speeds over a 10 minute period were determined using the different particle tracking methods, and are shown in figure 7. The results suggest that there was relatively good agreement between the manual and semi-automatic methods (differences in particle speeds only occurred at 5, 6 and 7 minutes), but there was a significant difference between the mean particle speed determined by the automatic and manual methods, and the automatic and semi-automatic tracking methods for most time-points. Hence this figure shows that the mean particle speed from the semi-automatic method is closer to the mean particle

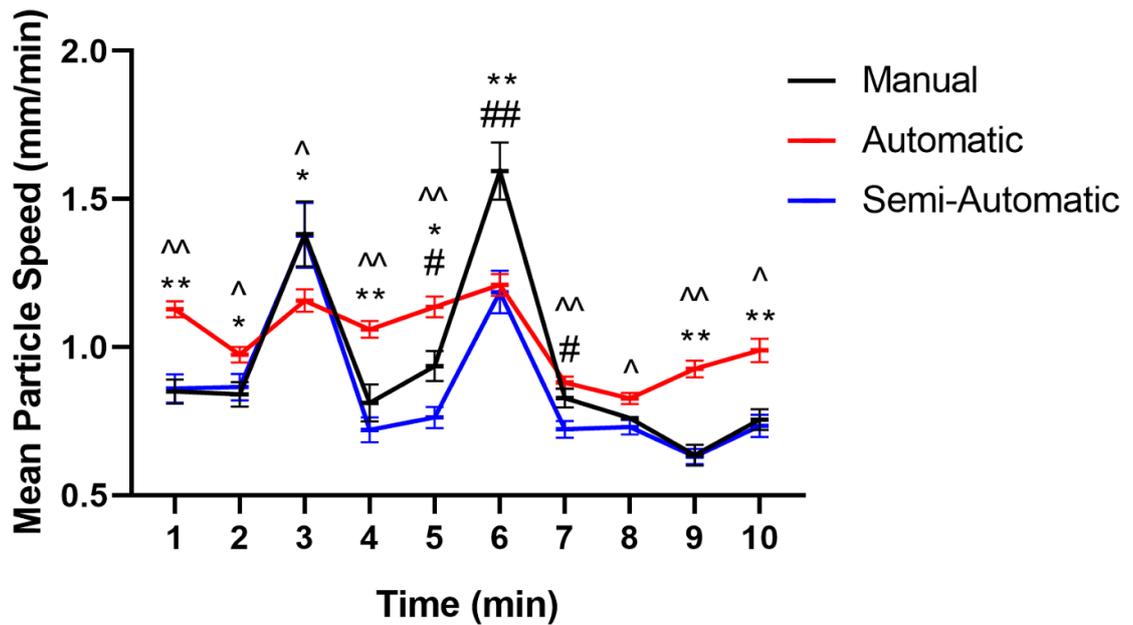


Figure 7. Average particle speed (Mean \pm SEM) for the different particle tracking methods. Stationary particles were not included in this analysis. Statistical analysis was conducted using a two-way repeated measures-ANOVA with Tukey corrections for multiple comparisons. * and ** indicate a difference between manual and automatic tracking methods with significance $p < 0.01$ and $p < 0.0001$ respectively. # and ## indicate a difference between manual and semi-automatic tracking methods with significance $p < 0.05$ and $p < 0.0001$ respectively. and ^ indicate a difference between semi-automatic and automatic tracking methods with significance $p < 0.05$ and $p < 0.0001$ respectively.

speed from the manual tracking method, and the automatic tracking method produces significantly different mean particle speeds. However, without knowing the ground-truth particle locations for the *in vivo* data we cannot say which particle tracking method produced the most accurate mean particle speed.

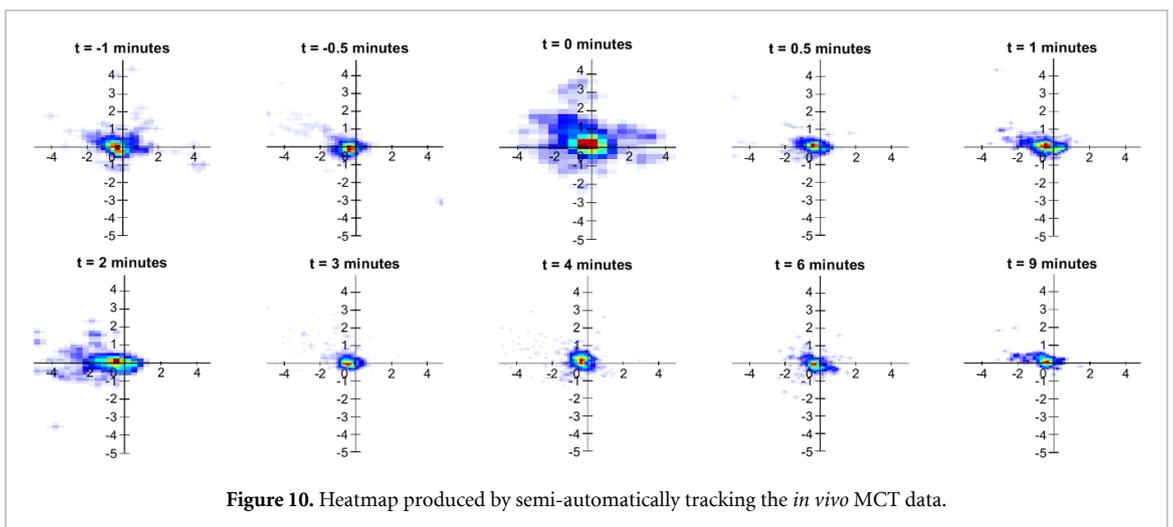
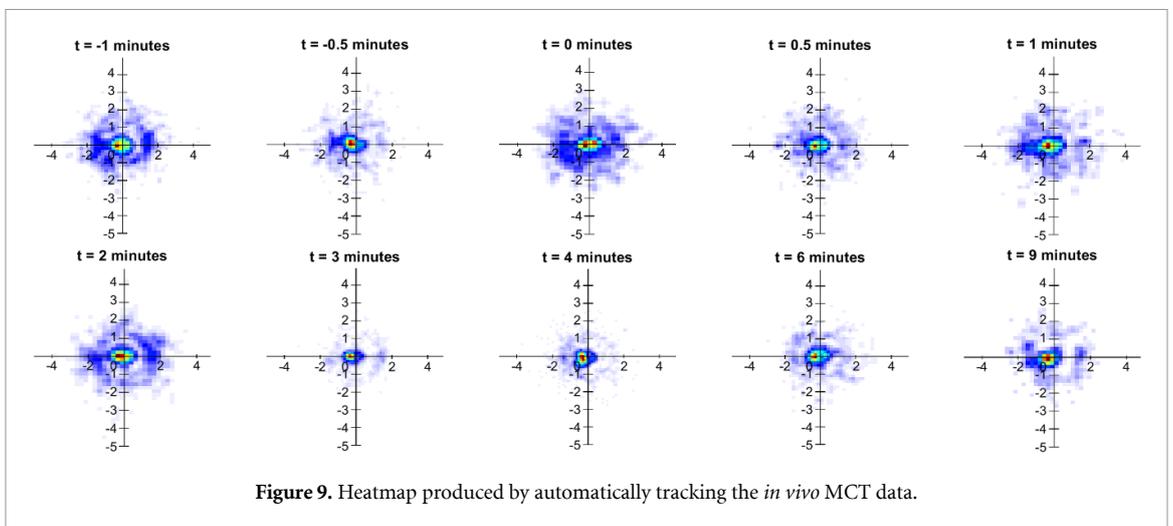
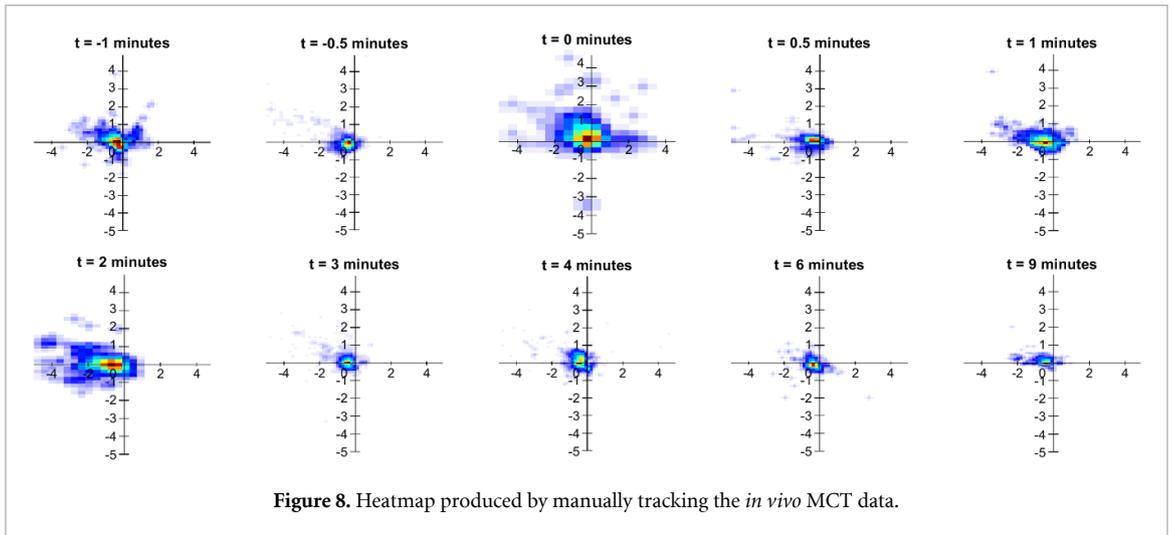
3.2.2.2. Variability in particle behaviours

Most particle-based MCT tracking techniques can only measure bulk particle motion and produce a mean MCT rate. However, one of the advantages of the S-PCXI approach is that the variability in individual particle behaviour can be interrogated. In previous studies we have visualised this by plotting the velocity of all of the individual particles, to create 2D histograms (Donnelley *et al* 2017, Donnelley *et al* 2019, Gardner *et al* 2019). These ‘heatmaps’, provide an additional level of detail about the MCT marker particle motion, showing the variability in the speed and direction of all the moving particles. The particle velocity heatmaps for the manual, automatic and semi-automatic particle tracking methods are shown in figures 8, 9, and 10. The heatmaps generated by the manual tracking method in figures 8 and 10 are similar, which echoes the data presented in figure 7. This shows that not only are the mean particle speeds similar, but as with the simulated data results, the particle direction detected using the semi-automatic method is also similar to the manual tracking method.

Figure 9 shows the effect that false particle tracking has on the heatmap results. The distribution of the heatmaps in 9 are different to those in 8 and 10 showing that for this dataset the fully-automatic particle detection method was not suitable for generating particle velocity heatmaps. Comparing figure 10 and figure 9 shows the effect that the false particle detection and tracking has on the accuracy of the distribution of the particle velocity heatmaps.

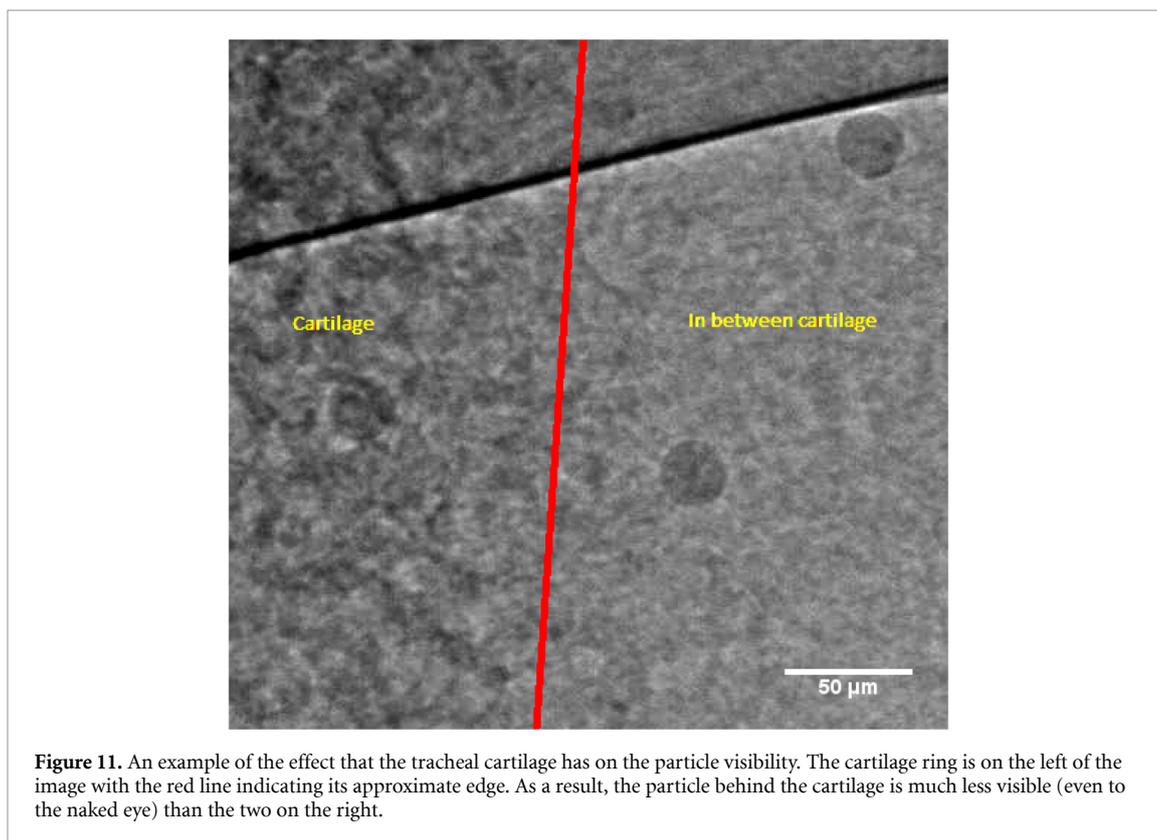
4. Discussion

When applied to the simulated data the automatic particle tracking method was able to track particles without any user input and was able to accurately detect the particle locations. However, the accuracy of the automatic particle tracking, and hence the estimation of particle velocities, was significantly lower than the manual and semi-automatic tracking methods. Conversely, whilst the manual tracking method was able to accurately link the particle locations between frames, the particle positional error was larger than the automatic tracking method, and a very large amount of time was required to analyse the data. The semi-automatic method proved to be a good compromise between the two particle tracking strategies, and



when using this method the particles could be tracked more accurately than with either the manual and automatic tracking methods, and required less time than the manual tracking method.

The automatic and semi-automatic methods presented in this paper were designed specifically for synchrotron PCXI images acquired to measure *in vivo* MCT speed and behaviour. The primary advantage of this method is that the spatial variation in individual MCT marker particle velocities can be measured, as opposed to other methods of monitoring MCT using MCT marker particles that measure the time taken for one or a group of particles to travel a known distance (Rogers *et al* 2018, Grubb *et al* 2016, Hua *et al* 2010).



This allows the heatmaps, shown in figures 8, 9 and 10, to be generated as well as the mean particle speed plots. Additionally, for these alternate MCT monitoring methods, the particles are delivered in groups to assist in visualising the particle positions. However, it is possible that the methods developed in this paper can be adapted to these alternate methods to assist in MCT measurement.

The results in table 2 show that by combining methods for particle detection, the FPR can be reduced without a significant decrease in the particle detection accuracy, but it does come at the cost of processing time. However, since current applications of the particle tracking algorithms are not applied in real-time, having a slightly longer processing time is not as big a disadvantage as a decrease in accuracy.

Comparing the data in table 1 to the data in table 2, the sensitivity of the particle detection algorithm is higher when applied to the simulated data, and the false negative and false positive rates are also lower when compared to the performance on the *in vivo* data. The differences between the sensitivity, false positive and false negative rates in table 2 and table 1 could also be due to the number of frames that were analysed. The results in table 2 were generated from analysis on 10 images, whereas the results in table 1 were generated from 500 images.

However, the differences are most likely due to a lack of anatomical landmarks in the simulated data images. Examples of features that can cause problems in accurate particle detection include the cartilage of the tracheal rings, and the airway surface itself as shown in figure 11. In the S-PCXI images the cartilage can create patterns that obscure the real particles, making them more difficult to detect. They can also create patterns that are falsely detected as particles, accounting for the differences in the sensitivity and FPRs in tables 1 and 2.

The results in figure 6 suggests that the automatic particle tracking method will underestimate the mean particle speed. Unsurprisingly, this suggests that for faster moving particles the automatic particle detection method struggles to link the particle locations into trajectories. However it could also be due to the automatic tracking algorithm linking particles that are close together, which would add more lower particle speed values, reducing the overall average particle speed.

The results of the mean particle speed measured from the *in vivo* images (figure 7) showed differences in measured particle speed between the different particle tracking methods at several time-points. The differences between the semi-automatic and manual particle tracking methods at time-points 5, 6 and 7 minutes are due to several different factors. The results of the simulated data in table 3 and figure 5 showed that there was a larger positional error for the manual particle detection than for the automatic particle

detection method. These differences in particle location could have contributed to the differences in the mean particle speeds at these time-points.

Another reason for this difference is the way in which a particle is classified as moving or stationary. For the manual particle linking this process is completely up to the user. For the semi-automatic particle linking the total particle displacement over the 15 seconds worth frames for that time-point is measured, and if that displacement exceeds a threshold then the particle path is shown to the user. Hence, whilst the final decision as to whether the particle is classified as moving or stationary is up to the user for the semi-automatic method, the user can also see which particles have moved a certain distance, which can help decide whether the particle is designated as moving or stationary. By classifying more particles as 'moving' in the semi-automatic method (especially slower moving particles), the mean particle speed is reduced, and is one of the reasons for the difference in the particle speeds between the manual and semi-automatic particle methods.

One of the advantages of the manual particle detection over the automatic particle detection is when the particles are grouped together. As mentioned previously, when the particles are grouped together, the cascade classifier has trouble identifying all of the particles in a group consistently. Whilst this could potentially be improved with further training, future work will need to look at new strategies for increasing the effectiveness of automatic particle detection algorithms for groups of particles.

The tendency of the particles to clump together, as shown in figure 3, is a recognised feature of the MCT system. The particles can clump together because of the physical nature of mucus production from goblet cells, and the presence of mucus tracts that form due to local ciliation patterns. Some regions of the airway have fewer cilia, so once the particles reach those regions they can become 'stuck', and can only be cleared with the addition of bulk fluid or by cough clearance. Particle clumping could also be enhanced by particle physicochemical properties and interactions with the airway surface, such as alterations in surface tension. Regardless of the cause, this clumping makes the accurate detection and tracking of individual particles substantially more difficult.

The primary limitation of the manual tracking method is the time taken to analyse the data. Whilst humans are good at detecting the particle motion, particularly for hard-to-detect particles and for complex particle trajectories, humans can perform poorly with monotonous and repetitive tasks taking large amounts of time. Hence the time taken to analyse the data only looked at the manual tracking time, and not the time taken for the automatic (and semi-automatic) tracking algorithms to train the particles detectors, and to run the particle detection and particle linking algorithms. This was because manual tracking required a researcher to be sitting at the computer actively processing the data, whereas the automatic algorithms could be run overnight or in the background whilst other tasks are being performed, since a dataset of 600 images takes approximately one hour to process. Hence it was decided that it was more important to reduce the amount of manual tracking time needed, rather than look at the total time required to process the data.

The results in table 3 show that the semi-automatic method was able to reduce the amount of manual tracking time required by 34%. However the semi-automatic method still required several hours of manual tracking time. As mentioned previously, the advantage of the semi-automatic particle tracking method is that it can track the trivial particle trajectories, leaving the more complex particle paths for a person to manually track. From the results in table 3 it appears that tracking these complex trajectories makes up a large proportion of the manual tracking time. These should be the focus of future work, as they are likely to have the greatest impact on reducing the amount of manual input.

The main difficulties with the automatic particle linking of the MCT marker particles is the relatively large time interval between the images. As mentioned previously, this is because images are taken in the relatively quiescent period between inspiratory movements to avoid the significant loss of image quality that respiration motion causes. Other studies we have performed suggest that allowing animals to breathe spontaneously, rather than using a ventilator to generate regular breathing, means the sensitive MCT processes are less perturbed. However, the absence of mechanical ventilation and the image-capture trigger signals available can add variation to the time difference between images because respiration rates normally vary substantially over time. Future experimental methods must look at developing ways to reduce the time interval between images—for example, by acquiring multiple images between breaths—whilst not compromising the realism and usefulness of the measured *in vivo* results.

Other applications of particle detection and tracking are set up in a way that the particle itself is a local maximum, and the background is a minimum pixel value. This property can then be integrated into the particle linking algorithm by estimating the likelihood the particle is a true particle based on the pixel values associated with that particle (Kim *et al* 2015, Newby *et al* 2018, Sbalzarini and Koumoutsakos 2005). However, the nature of PCXI and the MCT marker particle properties causes a unique image footprint for

the particles (as seen in the S-PCXI images presented in this paper), where the maximum pixel values occur at the particle edges—due to strong phase-contrast effects—and the pixel values inside the particles are similar to those surrounding the particle. As well as requiring a unique method to detect the particles, the tracking particle profile does not allow for simple integration into existing particle tracking algorithms, as there is no trivial property to estimate the likelihood that the particle is not a false detection. Future work should look at attempting to generate a score to estimate particle accuracy, or using a method such as a convolutional neural network, in which a score can be assigned for each predicted location. This score could then be integrated into the particle linking algorithms, increasing both the detection and tracking accuracy.

Given the low frame rate and the relatively low particle detection accuracy, a potential solution for improving the particle linking effectiveness could be to use a multiple hypothesis tracking (MHT) method. In an objective comparison of different particle tracking methods Chenouard *et al* (2014) suggested that multi-frame methods such as MHT are superior for particle tracking. This approach more closely resembles the approach that the human operator would use for manual tracking, in that the particle motions are observed over a series of consecutive frames prior to logging the manual tracking, as humans are better at detecting this type of motion than small differences between two frames. However for this particular application up to 200 particles could be present within a frame, so this method would generate an extremely large number of hypotheses. Typically to reduce the number of possible hypotheses when solving a MHT solution (and hence reduce the computational load), gating is used to reduce the hypotheses to only linking particles that are realistically close together (Kim *et al* 2015). However, as discussed previously, and shown in figure 3, the biological nature of the MCT marker particles on and inside the surface mucus is to group together over time, so although gating may reduce the number of potential hypotheses for the MHT tracking method, the computational load to track the number of potential particles over multiple frames would still likely be prohibitively large.

The automatic particle linking in this paper used a very simplistic particle estimation model to predict the particle location in the next frame. This simplistic model was chosen since the dynamics of the MCT marker particles are highly variable (Donnelley *et al* 2014b, Donnelley *et al* 2019, Rogers *et al* 2018) (see supplementary video 1 is available online at <https://stacks.iop.org/PMB/65/145012/mmedia> for a simulated image sequence and supplementary video 2 for the real image sequence), and we sought to minimise the effect that false particle linking and false particle detection had on the automatic tracking results. Future work will look at developing more advanced predictive models for particle movement and more intelligent predictive methods for the particle location. This is important because Chenouard *et al* noted that the accuracy of the prediction models can affect the particle tracking performance just as much as the tracking method itself (Chenouard *et al* 2014). Hence we propose that in order to improve the effectiveness of the particle prediction accuracy, more work is needed to better understand the *in vivo* dynamics of the MCT marker particles.

As mentioned in the introduction, the intended future application of this method is examining the effect of airway gene therapies on MCT as a step towards developing a cure for CF. The effects on MCT are likely to be small, and sample size calculations have shown that large cohorts of animals are needed in each group. Analysing the results from these large datasets is not feasible with the manual tracking method, as it would require many months of manual tracking. A semi-automated method opens up options for large studies. Additionally, it allows for a more thorough analysis of previously acquired datasets at more timepoints, allowing for greater time resolution of the results.

5. Conclusion

Automatic and semi-automatic particle tracking algorithms have been developed for the purpose of quantifying MCT behaviour by non-invasively tracking marker particle motion via synchrotron PCXI. The semi-automatic and automatic methods were tested on simulated and real world *in vivo* data, and compared with manual particle tracking. The semi-automatic tracking method reduced the time taken to analyse the MCT marker particle data by 34%, and is significantly more accurate than the automatic method developed. The strengths of the semi-automated method mean that in our laboratory it is now applied to the analysis of all synchrotron PCXI datasets. Additionally, the further development of particle tracking algorithms, as well as changes in the synchrotron PCXI image acquisition methods, will allow for continued development of fully-automatic MCT particle tracking methods in the future.

Acknowledgements

Experiments developed and performed on the BL20XU beamline at the SPring-8 Synchrotron under proposal numbers 2017A1301 and 2017B1386. Studies supported by NHMRC Project Grant GNT1079712, a Fay Fuller Foundation Discovery Grant, and the Cure4CF Foundation. MD was supported by a Robinson Research Institute Career Development Fellowship, and KM by a Veski VPRF Fellowship, then an ARC Future Fellowship FT180100374. KM completed this work with the support of the Technische Universität München Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under Grant Agreement No. 291763 and co-funded by the European Union. Authors were supported by the Australian Synchrotron International Synchrotron Access Program.

Data statement

The datasets analysed during the current study are available from the corresponding author on reasonable request.

ORCID iDs

Mark Gardner  <https://orcid.org/0000-0003-4340-2547>

David Parsons  <https://orcid.org/0000-0003-1746-3290>

Kaye Morgan  <https://orcid.org/0000-0001-6997-7977>

Alexandra McCarron  <https://orcid.org/0000-0002-6045-3998>

Patricia Cmielewski  <https://orcid.org/0000-0002-2236-9410>

Martin Donnelley  <https://orcid.org/0000-0002-5320-7756>

References

- Boucher R C 2019 Muco-obstructive lung diseases *N. Engl. J. Med.* **380** 1941–53
- Chenouard N, Smal I, de, Chaumont F, Maška M, Sbalzarini I F, Gong I F and Meijering E 2014 Objective comparison of particle tracking methods *Nat. Methods* **11** 281–9
- Cystic Fibrosis Foundation 2019. 2018 Cystic Fibrosis Foundation Patient Registry Highlights *Tech. Rep.* (Bethesda, MD: Cystic Fibrosis Foundation)
- Donnelley M, Morgan K, Farrow N, Siu K and Parsons D 2016 Non-invasive airway health measurement using synchrotron x-ray microscopy of high refractive index glass microbeads *AIP Conf. Proc.* **1696** 020011
- Donnelley M, Morgan K S, Awadalla M, Farrow N R, Hall C and Parsons D W 2017 High-resolution mucociliary transport measurement in live excised large animal trachea using synchrotron x-ray imaging *Respir. Res.* **18** 95
- Donnelley M, Morgan K S, Gradl R, Klein M, Hausermann D, Hall C and Parsons D W 2019 Live-pig-airway surface imaging and whole-pig CT at the Australian Synchrotron Imaging and Medical Beamline *J. Synchrotron Radiat.* **26** 175–83
- Donnelley M, Morgan K S, Siu K K W, Farrow N R, Stahr C S, Boucher R C and Parsons D W 2014a Non-invasive airway health assessment: synchrotron imaging reveals effects of rehydrating treatments on mucociliary transit in-vivo *Sci. Rep.* **4** 3689
- Donnelley M, Morgan K S, Siu K K W, Fouras A, Farrow N R and Carnibella R P et al 2014b Tracking extended mucociliary transport activity of individual deposited particles: longitudinal synchrotron X-ray imaging in live mice *J. Synchrotron Radiat.* **21** 768–73
- Endrizzi M 2018 X-ray phase-contrast imaging *Nucl. Instrum. Meth. A* **878** 88–98
- Gardner M, McCarron A, Morgan K, Parsons D and Donnelley M 2019 Particle coating alters mucociliary transit in excised rat trachea: a synchrotron x-ray imaging study *Sci. Rep.* **9** 10983
- Grubb B R, Livraghi-Butrico A, Rogers T D, Yin W, Button B and Ostrowski L E 2016 Reduced mucociliary clearance in old mice is associated with a decrease in Muc5b mucin *Am. J. Physiol. Lung Cell. Mol. Physiol.* **310** L860–7
- Hua X, Zeman K L, Zhou B, Hua Q, Senior B A, Tilley S L, Bennett W D 2010 Noninvasive real-time measurement of nasal mucociliary clearance in mice by pinhole gamma scintigraphy *J. Appl. Physiol.* **108** 189–96
- Jaqaman K, Loerke D, Mettlen M, Kuwata H, Grinstein S, Schmid S L, Danuser G et al 2008 Robust single-particle tracking in live-cell time-lapse sequences *Nat. Methods* **5** 695–702
- Jonker R and Volgenant A 1987 A shortest augmenting path algorithm for dense and sparse linear assignment problems *Computing* **38** 325–40
- Jung H-W, Lee S-H, Donnelley M, Parsons D and Lee I 2017. Automated detection of circular marker particles in synchrotron phase contrast X-ray images of live mouse nasal airways for mucociliary transit assessment *Expert Syst. Appl.* **73** 57–68
- Jung H-W, Lee S-H, Donnelley M, Parsons D, Stamatescu V and Lee I 2019. Multiple particle tracking in time-lapse synchrotron X-ray images using discriminative appearance and neighbouring topology learning *Pattern Recognit.* **93** 485–97
- Kim C, Li F, Ciptadi A and Rehg J M 2015 Multiple hypothesis tracking revisited 2015 *IEEE Int. Conf. on Computer Vision (Santiago, Chile, 7–13 December 2015)* pp 4696–704
- McIntyre C, Donnelley M, Rout-Pitt N and Parsons D 2018 Lobe-specific gene vector delivery to rat lungs using a miniature bronchoscope *Hum. Gene Ther. Methods* **29** 228–35
- Morgan K.S., Parsons D., Cmielewski P., McCarron A., Gradl R., Farrow N., Siu K., Takeuchi A., Suzuki Y., Uesugi K. and Uesugi M. 2020 Methods for dynamic synchrotron X-ray respiratory imaging in live animals *J. Synchrotron Radiat.* **27** (submitted)
- Munkres J 1957 Algorithms for the assignment and transportation problem *J. Soc. Ind. Appl. Math.* **5** 32–8
- Newby J M, Schaefer A M, Lee P T, Forest M G and Lai S K 2018 Convolutional neural networks automate detection for tracking of submicron-scale particles in 2D and 3D *Proc. Natl. Acad. Sci. USA.* **115** 9026–31
- Nixon M S and Aguado A S 2012 *Feature Extraction and Image Processing for Computer Vision* (New York: Academic)

- Ojala T, Pietikainen M and Maenpaa T 2002 Multiresolution gray-scale and rotation invariant texture classification with local binary patterns *IEEE Trans. Pattern Anal. Mach. Intell.* **24** 971–87
- Rogers T D, Ostrowski L E, Livraghi-Butrico A, Button B and Grubb B R 2018 Mucociliary clearance in mice measured by tracking trans-tracheal fluorescence of nasally aerosolized beads *Sci. Rep.* **8** 14744
- Sbalzarini I and Koumoutsakos P 2005 Feature point tracking and trajectory analysis for video imaging in cell biology *J. Struct. Biol.* **151** 182–95
- Viola P and Jones M J 2004 Robust real-time face detection *Int. J. Comput. Vis.* **57** 137–54