

Particle Tracking with Neighbourhood Similarities: A New Method for Super Resolution Ultrasound Imaging

Andrew Mobberley¹, Georgios Papageorgiou¹, Mairead Butler¹, Evangelos D. Kanoulas², Julian Keanie³, Daniel Good³, Kevin Gallagher³, Alan McNeil³, Vassilis Sboros¹ and Weiping Lu¹

¹*Institute of Biological Chemistry, Biophysics and Bioengineering, Heriot Watt University, Edinburgh, U.K.*

²*Janssen Pharmaceuticals R&D, High Wycombe, U.K.*

³*Western General Hospital, Edinburgh, U.K.*

Keywords: Super Resolution Ultrasound Imaging, Particle Tracking, Microbubbles.

Abstract: Single particle tracking (SPT) is a method for the observation of the motion of individual particles within a medium. It is broadly used to quantify the dynamics of particle flow, such as molecules/proteins in life sciences. In this paper, we will improve the performance of SPT by considering the local neighbourhood dynamical and structural information of a particle when it is tracked in a medium through consecutive frames, referred to as particle tracking with neighbourhood similarities (PTNS). This method is applied to track microbubbles in contrast enhanced ultrasound. We will test the method on synthetic data for method validation before applying to animal and human prostate data. We show that PTNS can make a significant improvement in the tracking performance in synthetic data, and in animal data it was able to accurately produce complex structures. In human prostate data, we find that by varying the control parameters we can inspect different behaviours of the tracks and from that understand the characteristics of the blood vessels they travel along.

1 INTRODUCTION

Contrast-enhanced ultrasound (CEUS) is an imaging modality used in hospitals to depict the circulation of organs within the body (Xu, 2009). Microbubbles (MBs) that act as contrast agents are injected into the patient intravenously before the imaging takes place to enhance the ultrasound image (Sboros, 2008). While CEUS has a higher image contrast compared to conventional ultrasound, its resolution is still limited by the fundamental wave diffraction limit, which is around 1mm for ultrasound waves (Couture et al., 2018). Such resolution limitations also occur in other imaging devices, such as optical microscopy due to light diffraction.

Several methods have recently been developed in super resolution that can surpass the fundamental diffraction limit, referred to as super resolution imaging. Photoactivated localisation microscopy (PALM) is one such technique that can localise fluorophores in optical images with spatial accuracy much better than the diffraction limit, thus achieving super resolution optical imaging many times below the diffraction limit (Betzig et al., 2006; Hess et al., 2006). This

work in optical microscopy has inspired a similar approach in ultrasound by localising MBs in CEUS data (Couture et al., 2018; Errico et al., 2015). Super resolution ultrasound imaging (SRUI) is considered a major development in the field since the discovery of ultrasound imaging.

The key to achieving SRUI is particle tracking, which comprises three steps: particle detection, localisation and linking. Detection of MBs can be achieved either by using particle probability images (Yang et al., 2012) or by using comparisons to a reference MB signal (Brown et al., 2019). Since they are much smaller than the imaging wavelength, they need to be localised, for example, by identifying the centroids of the MB signals (Couture et al., 2018; Errico et al., 2015; Yang et al., 2012; Brown et al., 2019). The motion of these MBs is then approximated by linking the localised positions between consecutive frames.

Many linking strategies have been developed for particle tracking. However, they have all been developed under the common rationale that the MBs with the shortest distance between subsequent frames are most likely to be linked. The most popular of these




Figure 1: Visualisation of different linking strategies that can be employed for MB tracking. The blue, red and yellow circles represent MBs in the first, second and third frame respectively. The dashed lines represent the links. 1a shows the nearest neighbour method. 1b shows the motion model method. 1c shows the ground truth links.

methods is the nearest neighbour method (NN) (Mazzaferri et al., 2014; Godinez et al., 2009). An illustration of how NN links MBs is depicted in Figure 1a. According to this method, MBs in one frame are linked to the nearest MB in the next frame, regardless of flow direction or whether the MBs share the same vessel. For sparse MB fields, this method works well (Kanoulas et al., 2019), capable of constructing the vessel networks using MB tracks, even in low signal-to-noise ratio environments. However, as MB density increases, the number of errors produced by NN also increases (Opacic et al., 2018). To improve the linking accuracy of NN, motion models are added by imposing certain motion restrictions between frames, which is visualised in Figure 1b. These motion models can be simple, such as linear flows (Opacic et al., 2018; Tang et al., 2020), or more complex ones, such as the interactive multiple model (IMM) (Genovesio et al., 2006) which accommodate many different types of particle motion and also incorporate the particle size and intensity (Yang et al., 2012). Motion model linking strategies typically perform better than those which use the nearest neighbour method (Chenouard et al., 2014), and have produced good results when applied to CEUS images (Opacic et al., 2018; Tang et al., 2020; Kanoulas, 2020).

The main objective of SRUI is to visualise the structure and dynamics of vessels at resolutions down to the near-microscopic scale (Opacic et al., 2018; Lin et al., 2016). To achieve this objective, the linking of MBs between different blood vessels should

be avoided. Incorrect links affect the quality of the super resolution vascular reconstruction, causing the images to appear noisy. Moreover, MB linking should strictly follow the vessel structure even when the vessels bend and the frame frequency is limited, as is the case for clinical scanners; using a straight-line link in curved vessels reduces the spatial resolution. Additionally, even with an appropriate motion model for MB linking, the first link in a track remains problematic as there is no velocity information available. NN in this situation could not only result in a wrong track for itself but also affect neighbouring tracks. Finally, the point spread function of ultrasound imaging varies spatially (Kanoulas et al., 2019), meaning that the topological features of MBs are not robust enough to be used in the IMM filter. These issues can only be addressed with further improvements of the linking approaches available in the literature.

In this paper, we propose a new method which we call Particle Tracking with Neighbourhood Similarities (PTNS) to tackle the aforementioned issues. This method is based on two observations in our recent study of ultrasound imaging. First, since MBs are confined in vessels, MB linking should be avoided when there is not sufficient and continuous MB density along its path. Second, MBs around a local neighbourhood are likely to belong to the same vessel and so should have similar directions and speeds. These observations are analogous to how vehicles travel on a road.

This paper is arranged as follows. In Section 2.1, we provide an overview of PTNS. Sections 2.2, 2.3

and 2.4 describes the method in more detail. In Section 3, we describe the experiments used for testing PTNS, and finally we conclude by discussing our findings.

2 MATERIALS AND METHODS

2.1 Particle Tracking with Neighbourhood Similarities

To incorporate the above mentioned observations into PTNS, we first measure the MB density to define criteria that each link must meet in order to be accepted. These criteria are called density assisted linking, which will prevent links from being made between adjacent vessels. Second, we use the collective speed and direction information within a local neighbourhood to guide the linking. This is called velocity assisted linking, which will assist the first link in a track.

PTNS is implemented by running the particle tracking process twice, referred to as a two run process, as depicted in Figure 2. In the first run (R1), shown in the left dashed box, MBs are detected using the particle probability (Yang et al., 2012). Linking is performed using the NN and motion model methods, where each potential link is assigned a cost within a matrix. The cost is based on the distance between MBs in adjacent frames. This cost matrix is then minimised to find the appropriate link for each MB (Yang et al., 2012). Density assisted linking is applied here to refine the tracks. In the second run (R2), shown in the right dashed box, we apply velocity assisted linking using the velocity information gathered in R1. This will further refine the linking after density assisted linking has been applied, particularly for the first link in a track. Moreover, we introduce a new maximum density seeking method to re-evaluate the links that were rejected by density assisted linking to find whether an alternate, bending link exists (instead of a straight line link). These three new linking strategies, i.e. density assisted linking, velocity assisted linking and maximum density seeking, will be described in more detail in Sections 2.2, 2.3, 2.4 respectively.

2.2 Microbubble Density Assisted Linking

Our linking strategy involves minimising the cost matrix \mathbf{C} . For the MBs identified in frame t , \mathbf{m}_i where $i = 1, \dots, M$ are the number of MBs, we predict their

positions in frame $t + 1$, \mathbf{m}'_i , by using the motion models. These are then compared to the MBs identified in frame $t + 1$, \mathbf{n}_j where $j = 1, \dots, N$ are the number of MBs in the next frame. The cost matrix is constructed as

$$\mathbf{C} = \begin{pmatrix} d(\mathbf{m}'_1 - \mathbf{n}_1)^2 & d(\mathbf{m}'_1 - \mathbf{n}_2)^2 & \dots & d(\mathbf{m}'_1 - \mathbf{n}_N)^2 \\ d(\mathbf{m}'_2 - \mathbf{n}_1)^2 & d(\mathbf{m}'_2 - \mathbf{n}_2)^2 & \dots & d(\mathbf{m}'_2 - \mathbf{n}_N)^2 \\ \vdots & \vdots & \ddots & \vdots \\ d(\mathbf{m}'_M - \mathbf{n}_1)^2 & d(\mathbf{m}'_M - \mathbf{n}_2)^2 & \dots & d(\mathbf{m}'_M - \mathbf{n}_N)^2 \end{pmatrix}, \quad (1)$$

where

$$d(\mathbf{m}'_i - \mathbf{n}_j)^2 = (m'_{i,x} - n_{j,x})^2 + (m'_{i,y} - n_{j,y})^2, \quad (2)$$

is the distance and $\mathbf{m}'_i = (m'_{i,x}, m'_{i,y})$ and $\mathbf{n}_j = (n_{j,x}, n_{j,y})$ are the coordinates of the MBs. Links are made by minimising each row of equation 1 and optimised so that each MB is linked to and by one MB in two consecutive frames.

As discussed in Section 1, density assisted linking is a key restriction to prevent links from being made between different vessels. For an image sequence under investigation, the MB density, $\rho(x, y)$, in pixel (x, y) is defined as the number of MBs localised in this pixel for all image frames. Since MBs travel within vessels, correct links between \mathbf{m}_i and \mathbf{n}_j should satisfy the following minimum MB density criteria, which in turn update \mathbf{C} . The first condition is,

$$\mathbf{C}(i, j) = \begin{cases} d(\mathbf{m}'_i - \mathbf{n}_j)^2 & \text{if } \bar{\rho} > c_\rho (\rho(\mathbf{m}_i) + \rho(\mathbf{n}_j)) \\ \infty & \text{otherwise} \end{cases}, \quad (3)$$

where $\bar{\rho}$ is the average MB density over all the pixels along the straight-line path between \mathbf{m}_i and \mathbf{n}_j . Equation 3 ensures that there is sufficient MB density along the trajectory between the start and end pixels, which will prevent links from being made between neighbouring vessels (Kanoulas, 2020). The hyperparameter c_ρ , defined as a percentage of $(\rho(\mathbf{m}_i) + \rho(\mathbf{n}_j))$, is a control parameter for accepting or rejecting a link.

The second density assisted linking condition is,

$$\mathbf{C}(i, j) = \begin{cases} d(\mathbf{m}'_i - \mathbf{n}_j)^2 & \text{if } \sigma_\rho < c_\sigma \bar{\rho} \\ \infty & \text{otherwise} \end{cases}, \quad (4)$$

where σ_ρ is the standard deviation of the MB density over all the pixels along the straight-line path between \mathbf{m}_i and \mathbf{n}_j and c_σ is a hyperparameter. Equation 4 ensures that the MB density does not fluctuate significantly along the path between \mathbf{m}_i and \mathbf{n}_j . The hyperparameter c_σ controls how large the fluctuation of the MB density can be before the link is rejected.

In summary, the density criteria given in equations 3 and 4 require that along the trajectory of a track, there must be a sufficient number and a smooth distribution of MBs, thereby avoiding jumps of MBs between vessels and improving the vessel reconstruction.

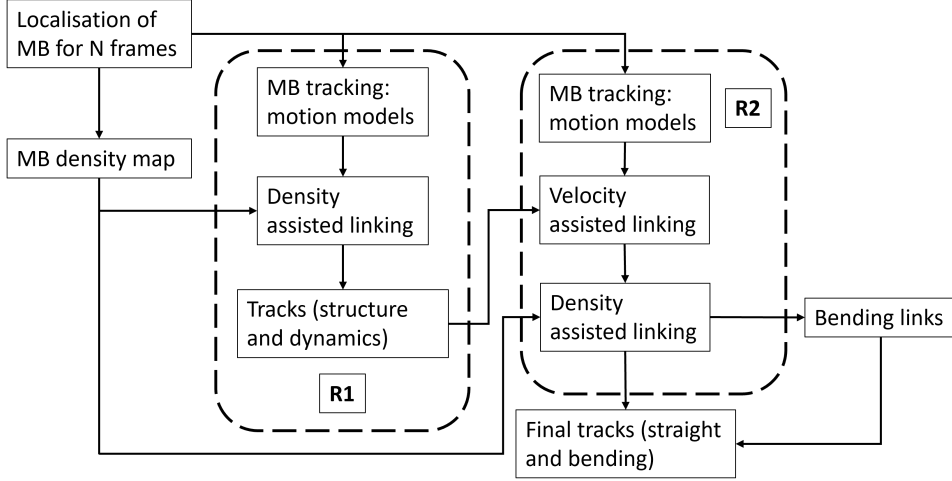


Figure 2: Flow chart of PTNS applied to CEUS data. The first run is shown in the dashed box labeled R1 and the second run is shown in the dashed box labeled R2.

2.3 Velocity Assisted Linking

As described earlier, we use motion models to better predict how MBs move within vessels, but these require velocity information from previous links. For the first link in a track, there is no velocity information and so the link can only be made using NN. We have developed a method to guide the links based on speed and direction information provided by the local neighbourhood: velocity assisted linking.

We implement this concept in the second run of PTNS (R2), which is depicted by the right dashed box in Figure 2. On the completion of the first run (R1), we can use all the MB tracks to generate their speed and direction maps, $v(x,y)$ and $\theta(x,y)$, which are the average speed and direction of the tracks passing through pixel (x,y) respectively. We can further calculate the standard deviation of the speed and direction of these tracks passing through each pixel, $\sigma_v(x,y)$ and $\sigma_\theta(x,y)$ respectively. These four measurements will then be incorporated into the linking in R2.

When a MB travels along a vessel, its speed and direction vary between neighbouring frames but should be within a certain continuity range. Here, we introduce two velocity assistance criteria, fixed and variable continuities respectively, suitable to different data sets under investigation. Which criterion is best applicable depends on the results obtained in R1. The fixed continuity criterion is best applicable when all the tracks in the entire image have fairly good continuity in both speed and direction, and their variations both locally and globally are small and well defined.

In this case, the cost matrix \mathbf{C} can be updated as

$$\mathbf{C}(i,j) = \begin{cases} d(\mathbf{m}'_i - \mathbf{n}_j)^2 & \text{if } \beta_1 v(\mathbf{m}_i) < v < \beta_2 v(\mathbf{m}_i) \\ & \text{and} \\ & \theta(\mathbf{m}_i) - \Theta < \theta < \theta(\mathbf{m}_i) + \Theta \\ \infty & \text{otherwise} \end{cases}, \quad (5)$$

where v and θ are the speed and direction of the link between MBs \mathbf{m}_i and \mathbf{n}_j , and $\mathbf{m}_i = (m_{i,x}, m_{i,y})$ are the pixel coordinates of MB \mathbf{m}_i in frame t . The coefficients β_1 and β_2 are the scaling parameters giving the fixed lower and upper boundaries of the speed respectively, $\beta_1 < 1$ and $\beta_2 > 1$. The angle Θ defines the fixed angle variation allowed around $\theta(\mathbf{m}_i)$. All three parameters can be estimated from R1.

The second criterion applies when the speed and direction continuity is not consistent, varying significantly across the image. In this case, the variable continuity criterion is more appropriate and \mathbf{C} is updated as

$$\mathbf{C}(i,j) = \begin{cases} d(\mathbf{m}'_i - \mathbf{n}_j)^2 & \text{if } v(\mathbf{m}_i) - \gamma\sigma_v(\mathbf{m}_i) < v < \\ & v(\mathbf{m}_i) + \gamma\sigma_v(\mathbf{m}_i) \\ & \text{and} \\ & \theta(\mathbf{m}_i) - \gamma\sigma_\theta(\mathbf{m}_i) < \theta < \\ & \theta(\mathbf{m}_i) + \gamma\sigma_\theta(\mathbf{m}_i) \\ \infty & \text{otherwise} \end{cases}, \quad (6)$$

where γ is a parameter that controls how narrow the continuity range is. A lower value of γ results in a more restrictive continuity range. As observed in equation 6, the restriction varies from pixel to pixel depending on the standard deviation of the speed and direction in that pixel.

The above two criteria given in equations 5 and 6 each have advantages and disadvantages and will be applicable to some data sets more than others. The




Figure 3: The potential links that can be made using R1 and R2. A MB in frame t is shown as a red circle, and the candidate MBs it could be linked to in frame $t + 1$ are shown as orange circles. The correct MB is shown with a green border. The gray arrows show the direction of the blood flow, and the dashed black lines show the potential links which can be made from the starting MB. There are more possible potential links in R1, 3a, than there are in R2, 3b. How velocity assisted linking restricts the potential links based on the speed and direction continuity is shown by the gray dashed lines in 3b.

former is homogeneous across the image but contains three free parameters to be adjusted, whereas the latter has a single parameter to set. As discussed earlier, which criterion is best suited with what parameter values is determined by the first run, which will be further discussed in the Section 3 when a specific data set is tested. Since the first link of any track can apply the speed and direction information in R2 from the other tracks in its neighbourhood, already available from R1, this link is no longer made using only the nearest neighbour method as in R1. Along with the motion models, all links produced by PTNS use the velocity information from the local neighbourhood, and thus all the tracks have a certain degree of continuity in both speed and direction, set by the control parameters. This process is visualised in Figure 3.

2.4 Maximum Density Seeking

When equations 3 and 4 are applied, a link previously accepted by the motion models may be accepted or rejected. If a link is accepted, then it can be used to form part of a track. If a link is rejected, it could genuinely be wrong, or it may be correct but the straight line path may be wrong as the vessel bends significantly at this location. For the latter case, an alternate, non-straight trajectory is more realistic. We propose a new maximum density seeking method to recover such a path after the straight-line link has been rejected.

The method finds a path with maximum MB density between two consecutive frames, which is referred to as bending links in Figure 2.

As shown in Figure 4a, after the straight-line link from m_i to n_j is rejected by the path density criteria, a pixel among the 8 neighbours of m_i , which has a MB density above a certain threshold and has the shortest distance to n_j compared to the other pixels, is chosen as the next stage in the linking process. This pixel is labeled as (x_1, y_1) in Figure 4b. This process is repeated by sliding the neighbourhood window from m_i to (x_1, y_1) as shown in Figure 4c. Continuing the process leads to successive link points (x_3, y_3) , (x_4, y_4) , ..., (x_N, y_N) , until n_j , forming the bending link from m_i to n_j shown in Figure 4d. If (x_N, y_N) cannot reach n_j , the process fails and no alternative path is possible.

3 RESULTS

To test the new linking strategies, we will apply PTNS to three different data sets: synthetic, animal and human prostate CEUS images. These data sets will allow us to quantitatively and qualitatively assess the performance of PTNS under various different conditions, some of which have distinct and known characteristics that can be compared to the structural and dynamical features recovered by PTNS.

3.1 Synthetic Data Test

A synthetic flow model is first generated to simulate blood flow in a vascular network. Point-like particles are then injected into the network, moving along the




Figure 4: The steps of the maximum density method. Pixels with high MB density are shown in red, orange and yellow. Pixels with low MB density are shown in green, and pixels with no MBs detected are shown in blue. The black circle and diamond show the starting and ending pixel, m_i and n_j , respectively. 4a: the straight-line link between m_i and n_j is rejected because it violates the density criteria. 4b: the neighbourhood around m_i is inspected and the nearest pixel to n_j which is above the pixel threshold is labeled (x_1, y_1) . 4c: the process is repeated by inspecting the neighbourhood around (x_1, y_1) to find (x_2, y_2) . 4d: The process continues, finding (x_3, y_3) , (x_4, y_4) and (x_5, y_5) . The bending link is then fitted from m_i to n_j through the pixels of highest density (x_1, y_1) to (x_5, y_5) .

flow. These particles are then blurred using a variable point spread function that mimics MBs observed in real in vivo CEUS data. White Gaussian noise is finally added to each frame. This leads to a realistic CEUS data set of MBs travelling with variable speeds along a pre-designed network structure, each of which is tagged with an ID number (Kanoulas et al., 2019) to be used as the ground truth to compare with the tracking results of PTNS.

To evaluate the performance of PTNS, we compare the speed maps of MBs obtained by different particle tracking methods with the ground truth, which is shown in Figure 5. In particular, we focus on the central region of the network where the MB density is low, and the vessels are thin and close together, which are all challenging conditions for MB tracking. The speed map registers the average speed of all tracks that pass through each pixel in the image sequence, which allows us to construct the vascular dynamics and structure. The actual speed map (ground truth) is shown in Figure 5a.

We first apply NN combined with the motion model to the synthetic data. As shown in Figure 5b, the overall structures are present, but those in the top and bottom of the image are weak and some of them

are missing. Moreover the thin and densely packed vessels in the centre are not well resolved. This is because incorrect links between neighbouring vessels are more likely to occur with NN.

We now apply PTNS to the synthetic data. The results for R1 are shown in Figure 5c, which has shown a significant improvement in the recovery of the overall network structure in comparison to the ground truth. In particular, the thin, densely packed vessels in the middle region are mostly recovered. This must be attributable to the density assisted linking as it is the main difference to NN. However, there are still some noticeable wrong links in R1, both in terms of speed and direction. The results of R1 show that the velocity continuity is mostly uniform across the network, and so we can apply the velocity assisted linking model equation 5 to perform the second run.

As shown in Figure 5d, most of the incorrect links in the middle region in R1 are removed in R2, based on the velocity information gathered in R1. We further apply the maximum density seeking method to include the bending links, the results of which are shown in Figure 5e. The main visual difference between Figures 5d and 5e is that the addition of bending links appears to make the speed map slightly nois-




Figure 5: Speed maps of the synthetic vessel network. All are in units of mm/s. 5a: Speed map produced from the ground truth. The region of interest is enclosed within the white dashed box. 5b: Speed map produced from the tracks generated by the NN and motion model method. 5c: Speed map produced from the tracks generated using PTNS in R1. 5d: Speed map produced from the tracks generated using PTNS in R2. Only the straight-line links are plotted. 5e: Speed map produced from the tracks generated using PTNS in R2. Both the straight and the bending links are plotted. $c_p = 0.15$ and $c_\sigma = 1.3$ and a fixed continuity range is used here, with $\beta_1 = 0.5$, $\beta_2 = 1.45$ and $\Theta = 45^\circ$ being used in equation 5.

ier. To assess the full impact of the bending links, we perform a quantitative analysis of the links, which can be seen in Table 1.

Table 1 shows the number of true positive links that a linking method produced (TP) and the total number of links produced by the method (TL). The total number of links in the ground truth (GT) is given in the caption. From these, we calculate two statistics, precision and Jaccard index, to quantitatively measure the performance of different linking methods. From left to right in the table: NN plus motion model method, PTNS R1, PTNS R2 with straight-line linking, PTNS R2 with bending linking, PTNS R2 with combined straight-line and bending linking.

We can see that the precision and the Jaccard index improve as we apply the different elements of PTNS (columns 2-5 in Table 1). The straight-line links produced by R1 show high precision, and this is improved in R2. However, as seen by the Jaccard index, only $\sim 33\%$ of the total links in the ground truth are produced. This inefficient use of the data justifies the use of the maximum density seeking method to recover links that have been missed. Combining the straight-line and bending links in R2 greatly improves the Jaccard index, recovering $\sim 50\%$ of ground truth

links. There is a small decrease in the precision due to the lower precision of the bending links in comparison to the straight-line links, which results in the noise seen in Figure 5e.

3.2 Animal Data Test

We now proceed to test PTNS on an animal data set. The data was acquired from a sheep ovary using a 1.2mL bolus injection of SonoVue (Bracco, Geneva, Switzerland) contrast agent. We collected 936 CEUS frames with a frame rate of 13 Hz. Here, the ground truth particle locations are no longer available. In vivo data may have location, architecture and vessel dimensions available through other techniques. Our previous work showed that the best ground truth for in vivo animal data may be acquired through optical coherence tomography (OCT) (Kanoulas et al., 2019), where the acquisition is live and vessel features are preserved and can be directly compared with the SRUI output. Figure 6 provides the output of PTNS as described in Figure 2. The original OCT image is in Figure 5b in (Kanoulas et al., 2019).

The MBs in each CEUS frame have their positions localised and collected into a MB number map,

Table 1: Validation results for the synthetic data. TP: number of true positive links produced. TL: total number of links produced. GT: total number of links in the ground truth. For this synthetic data set, $GT = 20040$.

	NN+MM	R1	R2, S links	R2, B links	R2, S+B links
TP	5800	6008	6856	3531	10387
TL	6820	6728	7109	3881	10990
Precision	0.8504	0.8930	0.9644	0.9098	0.9451
Jaccard (%)	27.54	28.94	33.79	17.32	50.32




Figure 6: Tracking results for the sheep dataset. 6a: MB number map. The white arrows indicate the three important structures in this data set. 6b: Speed map for NN tracking. 6c: Speed map for the links produced by R1. 6d: Speed map for the straight and bending links produced by R2. All speed maps are in units of mm/s. $c_p = 0.5$ and $c_\sigma = 1$ and a variable continuity range was used here, with $\gamma = 1.5$ being used in equation 6.

where each pixel (x, y) in the MB number map contains the total number of MBs localised in that pixel for all image frames, as previously described in Section 2.2. This result is shown in Figure 6a. Within the MB number map, there are three distinct features that can clearly be seen. These are labeled 1, 2 and 3 in Figure 6a and they are as follows: a horseshoe shaped structure, a curved structure, and a thin vessel. Reproduction of these three features will be used in place of the ground truth for this data set to assess the performance of PTNS. Just as with the synthetic data, we use speed maps to show the dynamics and structure, which will allow us to clearly visualise the results of each tracking method.

We first apply NN to the sheep data. From Figure

6b, it can clearly be seen that feature 1 is not reproduced well. The horseshoe shape is heavily polluted, which shows that there are many links being made across the gap in the middle of the structure. Feature 2 is mostly well reproduced, but there are links connecting it to feature 1 on the left hand side which are incorrect. Of the three features, feature 3 is the best reproduced. In this region, the MB concentration is low, the speed is low and the structure is thin and well isolated, which are conditions that NN performs well in.

Next we apply PTNS to the sheep data. The speed map for R1 is shown in Figure 6c, and the speed map for R2 combining straight-line and bending links is shown in Figure 6d. Features 1 and 2 are better re-

produced using PTNS than they are using NN. The empty region in the middle of the horseshoe shape in feature 1 is preserved in Figures 6c and 6d, and there are no longer any incorrect links connecting the left of feature 2 to feature 1. However, straight-line links are not enough to preserve the connection between features 1 and 2 on the right side of feature 2. Combining straight-line links with bending links in R2 reproduces this connection. It also further improves the structure of features 1 and 2, which can be seen in Figure 6d.

The thin vessel in feature 3 is not well reproduced in Figure 6c. There is a slight improvement in Figure 6d, though there are still parts missing. Because the MB density is sparse in this region, it is more likely for links to be rejected due to density assisted linking. Because of this sparsity of MBs, there isn't a clear path of highest MB density for the maximum density seeking method to follow, and thus it fails to find alternative, non-straight-line paths. As a result, this shows that the maximum density seeking method is not suitable for data sets with a very low MB concentration.

3.3 Human Prostate Data

Finally, we apply PTNS to human prostate data. CEUS data was collected in the Western General Hospital in Edinburgh from patients scheduled for radical prostatectomy with full ethical approval. An infusion of Lumivity[®] (Lantheus, Billerica, USA) contrast agent was administered to the patient with a roughly constant flow rate, and data was collected for 3-4 minutes at a frame rate of 10 Hz using an iU22 Philips scanner with C10-3v transducer. The transducer was secured in place, scanning a region suspected to contain cancer from prior MRI assessment. The ground truth here refers to the structure and location of the tumour as determined by histopathological evaluation. In comparison to the sheep data, this no longer specifies the vessel dimensions and architecture but rather the approximate topography of the tumour, though the results may not perfectly match the histopathology given the orientation of the ultrasound frame. Figure 7 shows the results of PTNS when applied to the prostate data.

Figure 7a is the B-mode prostate image of the patient under investigation, where the prostate border is circled with a yellow line. The histopathology results are shown in Figure 7b. Here the cancer is encircled by the black line. This was then matched onto SRUI images, taking into account the shape distortion due to the ultrasound scanner. We now apply PTNS to track the MBs in the CEUS data, beginning with R1. The track number map for R1, which

is the number of tracks passing through each pixel, is shown in Figure 7c while the corresponding speed map is shown in Figure 7d. The cancer region has been identified by our pathology team and is encircled by the blue dashed line in Figure 7c. As seen, the cancer region has both a high track number and a high speed. We note that high track number and high speed is also found in healthy areas, noticeably in the central region. However, there are no particular spatial structures in the cancer region compared to those in healthy areas.

We further apply R2 of PTNS, with a variable continuity range as defined by equation 6. Results with different continuity levels have been investigated by adjusting the control parameter γ . Figures 7e and 7f are the track number map and the speed map for R2 with $\gamma = 1.5$. Comparing Figure 7e with Figure 7c, it can clearly be seen that the total number of tracks in the prostate is reduced in R2, which also shows clearer structures due to the imposed continuity restrictions. Similar features due to continuity restrictions are also observed when comparing Figure 7f with Figure 7d.

From the literature, it is known that the vasculature of prostate cancer is different than that of healthy tissue (Forster et al., 2017; Alizadeh et al., 2013; Vaupel and Kelleher, 2012). Cancer blood vessels are typically thicker than healthy blood vessels (Forster et al., 2017; Müller et al., 2008), with increasingly tortuous vessels (Alizadeh et al., 2013) and a heterogeneous blood flow (Vaupel and Kelleher, 2012; Jochumsen et al., 2020). As such, we expect some discrepancies in the structures and dynamics between the cancer region and healthy region when the continuity range varies. In order to quantify this, a comparison between the number of tracks in the cancer and healthy regions at different levels of continuity restrictiveness is investigated. Table 2 compares the track number in the cancer and healthy regions of the patient for R1, and R2 with two different γ values.

As seen in Table 2, while both the cancer and healthy region shows a decrease in the number of tracks as the continuity range becomes more restrictive, the cancer region shows a greater reduction in track number; 70% compared to 75% for $\gamma = 1.5$, and 25% compared to 34% for $\gamma = 1$. The larger reduction in track number in the cancer region shows that this region lacks motion continuity more than the healthy region. This in turn indicates that the vessels and corresponding blood flow in cancer regions behave more chaotically in structure and dynamics respectively. This finding is in line with literature observations of prostate cancer blood vessels (Alizadeh et al., 2013; Vaupel and Kelleher, 2012; Jochumsen




Figure 7: PTNS results for the human prostate data set. 7a shows a B-mode image with the prostate enclosed by the yellow line. 7b shows the results of the histopathology. 7c shows the track number map for R1. The area enclosed by the blue dashed line is the cancer region. 7d shows the speed map for R1, in units of mm/s. 7e shows the track number map for R2. 7f shows the speed map for R2, in units of mm/s. $c_p = 0.15$ and $c_\sigma = 1.3$ and a variable continuity range was used here, with $\gamma = 1.5$ being used in equation 6.

et al., 2020). These findings show that PTNS can potentially be a useful tool to identify distinctive structural and dynamical features in cancer areas to assist in the diagnosis of prostate cancer.

4 CONCLUSIONS

In this paper, we proposed a new method which we call Particle Tracking with Neighbourhood Similarities which seeks to address the issues of incorrect

Table 2: Table showing how the number of tracks generated for the human prostate data set reduce as the strictness of the continuity range increases, presented as the raw number of tracks and also as percentage of the tracks generated in R1. The tracks in the cancer and healthy regions have been separated to show the difference in behaviour.

	Cancer Tracks	% of R1	Healthy Tracks	% of R1
R1	10386	100	46832	100
R2, $\gamma = 1.5$	7309	70.37	35379	75.54
R2, $\gamma = 1$	2569	24.74	15717	33.56

linking between vessels and the first link in a track identified in current particle linking approaches and methods when applied to CEUS images. In synthetic data, we have shown that PTNS improves the performance of particle tracking. PTNS uses the data more efficiently as a result of increasing the precision and Jaccard index. In animal data, we have demonstrated that PTNS can construct complex structures of in vivo data, particularly in regions of high MB density.

Finally, we have applied PTNS to human prostate data to investigate the structure and dynamics of the vasculature for a patient with prostate cancer. Because PTNS is a two run process, the second run can be used to probe different levels of motion continuity by varying the control parameters, which can provide the means to distinguish different blood flow dynamics between cancer and healthy tissue. Using this approach, we observe that the cancer region exhibited a more chaotic vascular structure, resulting in a larger decrease in track number when the velocity continuity was made more restrictive. This is consistent with the current understanding of prostate cancer. More broadly, this result shows that super resolution ultrasound imaging of prostate cancer can potentially be developed into tools for diagnosis and focal therapy.

ACKNOWLEDGEMENTS

We would like to thank Lantheus Medical Imaging, Inc. for providing the Luminity[®] contrast agent as part of their Research Grants program. We would also like to thank the Western General Hospital Edinburgh theatre team for their time and for allowing us to use their facilities.

REFERENCES

Alizadeh, F., Hadi, M., Khorrami, M., Yazdani, M., Joozdani, R., Tadayyon, F., and Mellat, M. (2013). Prostate cancer: Relationship between vascular diameter, shape and density and gleason score in needle biopsy specimens. *Advanced Biomedical Research*, 2(1).

Betzig, E., Patterson, G., Sougrat, R., Lindwasser, O.,

Olenych, S., Bonifacino, J., Davidson, M., Lippincott-Schwartz, J., and Hess, H. (2006). Imaging intracellular fluorescent proteins at nanometer resolution. *Science*, 313(5793):16–42–1645.

Brown, J., Christensen-Jeffries, K., Harput, S., Zhang, G., Zhu, J., Dunsby, C., Tang, M., and Eckersley, R. (2019). Investigation of microbubble detection methods for super-resolution imaging of microvasculature. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 66(4):676–691.

Chenouard, N., Smal, I., de Chaumont, F., Maška, M., Sbalzarini, I., Gong, Y., Cardinale, J., Carthel, C., Coraluppi, S., Winter, M., Cohen, A., Godinez, W., Rohr, K., Kalaidzidis, Y., Liang, L., Duncan, J., Shen, H., Xu, Y., Magnusson, K., Jaldén, J., Blau, H., Paul-Gilloteaux, P., Roudot, P., Kervrann, C., Waharte, F., Tinevez, J., Shorte, S., Willemse, J., Celler, K., van Wezel, G., Dan, H., Tsai, Y., de Solórzano, C., Olivo-Marin, J., and Meijering, E. (2014). Objective comparison of particle tracking methods. *Nature Methods*, 11(3):281–289.

Couture, O., Hingot, V., Heiles, B., Muleki-Seya, P., and Tanter, M. (2018). Ultrasound localisation microscopy and super resolution: a state of the art. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 65(8):1304–1320.

Errico, C., Pierre, J., Pezet, S., Desailly, Y., Lenkei, Z., Couture, O., and Tanter, M. (2015). Ultrafast ultrasound localisation microscopy for deep super-resolution vascular imaging. *Nature*, 527(7579):499–502.

Forster, J., Harriss-Philips, W., Douglass, M., and Bezak, E. (2017). A review of the development of tumor vasculature and its effects on the tumor microenvironment. *Hypoxia*, 5:21–32.

Genovesio, A., Liedl, T., Emiliani, V., Parak, W., Coppey-Moisand, M., and Olivo-Marin, J. (2006). Multiple particle tracking in 3-d+t microscopy: method and application to the tracking of endocytosed quantum dots. *IEEE Transactions on Image Processing*, 15(5):1062–1070.

Godinez, W., Lampe, M., Wörz, S., Müller, B., Eils, R., and Rohr, K. (2009). Deterministic and probabilistic approaches for tracking virus particles in time-lapse fluorescence microscopy image sequences. *Medical Image Analysis*, 13(2):325–342.

Hess, S., Girirajan, T., and Mason, M. (2006). Ultra-high resolution imaging by fluorescence photoactivation localisation microscopy. *Biophysical Journal*, 91(11):4258–4272.

Jochumsen, M., Sørensen, J., Pederson, B., Nyengaard, J., Krag, S., Frøkiær, J., Borre, M., Bouchelouche, K.,

- and Tolbod, L. (2020). Tumour blood flow for prediction of human prostate cancer aggressiveness: a study with rubidium-82 pet, mri and na+/k+-atpase-density. *European Journal of Nuclear Medicine and Molecular Imaging*, 48(2):532–542.
- Kanoulas, E. (2020). *Methodology for a Super-Resolution Contrast-Enhanced US 2D-Imaging for Clinical Diagnosis*. PhD thesis, Heriot Watt University.
- Kanoulas, E., Butler, M., Rowley, C., Voulgaridou, V., Diamantis, K., Duncan, W., McNeil, A., Averkiou, M., Wijkstra, H., Mischi, M., Wilson, R., Lu, W., and Sboros, V. (2019). Super-resolution contrast-enhanced ultrasound methodology for the identification of in vivo vascular dynamics in 2d. *Investigative Radiology*, 54(8):500–516.
- Lin, F., Rojas, J., and Dayton, P. (2016). Super resolution contrast ultrasound imaging: Analysis of imaging resolution and application to imaging tumor angiogenesis. In *2016 IEEE International Ultrasonics Symposium (IUS)*.
- Mazzaferri, J., Roy, J., Lefrancois, S., and S., C. (2014). Adaptive settings for the nearest neighbour particle tracking algorithm. *Bioinformatics*, 31(8):1279–1285.
- Müller, B., Lang, S., Dominietto, M., Rudin, M., Schulz, G., Deyhle, H., Germann, M., Pfeiffer, F., David, C., and Weitkamp, T. (2008). High-resolution tomographic imaging of microvessels. *Developments in X-Ray Tomography VI*.
- Opacic, T., Dencks, S., Theek, B., Piepenbrock, M., Ackermann, D., Rix, A., Lammers, T., Stickeler, E., Delorme, S., Schmitz, G., and Kiessling, F. (2018). Motion model ultrasound localization microscopy for preclinical and clinical multiparametric tumor characterization. *Nature Communications*, 9(1).
- Sboros, V. (2008). Response of contrast agents to ultrasound. *Advanced Drug Delivery Reviews*, 60(10):1117–1136.
- Tang, S., Song, P., Trzasko, J., Lowerison, M., Huang, C., Gong, P., Lok, U., Manduca, A., and Chen, S. (2020). Kalman filter-based microbubble tracking for robust super-resolution ultrasound microvessel imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(9):1738–1751.
- Vaupel, P. and Kelleher, D. (2012). Blood flow and oxygenation status of prostate cancers. *Advances in Experimental Medicine and Biology*, pages 299–305.
- Xu, H. (2009). Contrast-enhanced ultrasound: the evolving applications. *World Journal of Radiology*, 1(1):15–24.
- Yang, L., Qiu, Z., Greenaway, A., and Lu, W. (2012). A new framework for particle detection in low snr fluorescence live cell images and its application for improved particle tracking. *IEEE Transactions on Biomedical Engineering*, 59(7):2040–2050.