Fast automatic bone surface segmentation in ultrasound images without machine learning *

Shihfan Jack Tu[1][0000−0002−7186−9745], Jules Morel[1], Minsi Chen[2][0000−0001−6628−1421], and Stephen J. Mellon[1][0000−0002−6375−6839]

[1] Oxford Orthopaedic Engineering Centre, Nuffield Department of Orthopaedics, Rheumatology and Musculoskeletal Sciences, University of Oxford, Oxford, UK
{jack.tu,stephen.mellon}@ndorms.ox.ac.uk
https://www.ndorms.ox.ac.uk/research-groups/oxford-orthopaedic-engineering-centre
[2] Department of Computer Science, University of Huddersfield, Huddersfield, UK
M.Chen@hud.ac.uk

Abstract. Reconstructing 3D bone images with 2D clinical ultrasound image is one of the primary developmental trends of computer-assisted orthopaedic surgery procedures, and real-time bone segmentation is required for such development. We previously presented a dynamic programming method with local phase tensor extraction for bone structure segmentation that could process one ultrasound frame with a true positive ratio of 71% in approximately 1 second. The present study aimed to reduce the segmentation time to enable real-time computational capacity for clinical application developments. A simplified bone probability algorithm was optimised by systematically identifying and removing the components which cost most computing resources. The segmentation results produced by the bone probability method were compared to the local phase method, and manual segmentation carried out by clinical experts. The proposed method had higher recall metric (0.67) than the local phase method (0.61), while the computational time is reduced to 0.02 seconds per image. However, the bone probability method did not perform as well as the local phase method in specificity and precision metrics. In conclusion, the simplified version of the segmentation algorithm improved computational speed and promised an advantage in further real time application developments, but additional functions that can improve accuracy and further extensive validations are still required before further clinical application developments.

Keywords: bone segmentation · ultrasound imaging · bone probability map

* This work was supported by Wellcome Trust ISSF/University of Oxford MLSTF. The authors would like to thank Prof Irina Voiculescu and Mr Ziyang Wang for supporting and discussing the machine learning aspect. The authors also acknowledge the reviewers for the constructive feedback.

1 Introduction

Reliable automatic segmentation of bone in ultrasound images has a variety of potential applications within orthopaedics. Paired with a robust registration algorithm [7] and an external localiser, auto-segmentation can be used to accurately locate and track bony anatomy in 3D space.

The motion analysis with ultrasound system (CAT&MAUS) used in our work is composed of 2D B-mode ultrasound (US) scanners and an optoelectronic motion analysis system to directly locate and track under-skin bony landmarks for the purposes of measuring kinematics at joints. It has been previously applied to track patella motion during knee flexion [10] and the greater trochanter during gait [6, 8].

Although several methods exist, reliably and quickly segmenting the bone structure in US images remains a crucial challenge for the development of clinical systems. For example, US bone scans do not provide complete anatomical boundaries and shape, and the bone structures give a strong acoustic response and provide a shadowing effect below them [4]. An automatic bone segmentation algorithm using local phase features and dynamic programming [5] was developed to resolve the issue of strong acoustic response anatomical structures, such as fascia and tendon, being incorrectly recognized as a bone. This method reduced the need for manual segmentation and can be integrated with a motion capture system [9]. However, some limitations still need to be conquered to achieve the ultimate goal, which is to combine automatically segmented bone contours from 2D freehand US scans with positional information obtained from the probe transducer and to transform the data into a 3D bone model at a real-time pace. Although the automatic algorithm's outcome quality has the potential to enable this goal, the most critical issue to resolve is to reduce the long processing time.

Ultrasound image processing methods based on machine learning have been proposed for bone surfaces segmentation from ultrasound data. For example, deep learning network architecture, termed U-net and convolutional neural network [11, 12, 15, 14, 1]. A well-trained machine learning model could help with the bone model reconstruction's processing speed and promise wider applications for clinical usage. A previous study reported that after 6 hours of training, the machine learning programme based on the convolutional neural network (CNN) has the potential to segment one image in 52 milliseconds [1]. Although the speed seems promising, this method still requires a long local phase extraction procedure. This is a limitation that our previous method shares [5]. The local phase extraction was the most time-consuming processing step in the procedure. Considering the success of the machine learning methods is dependent on a number of training images, anatomical variation present in the training data, and the quality of the image. Manual labelling is still required in training data preparation. The proposed method in this paper can be used to reduce or remove this requirement.

The current study proposes methods to reduce the computation time to segment and register bone surface from specific ultrasound images, which contains

a high-intensity bone response profile, corresponding to the bone surface, followed by shadow region. Our target was to perform segmentation for one frame within 0.05 seconds in order to treat more than 20 frames per second, which is the common acquisition sampling rate of most US probes. Our ultimate developmental focus is to improve segmentation speed for use as a step within rapid reconstruction of bone position in 3D via registration. The secondary aim of the current study was to explore a fast non-machine-learning segmentation method that could offer acceptable accuracy and precision results with minimum human input.

2 Methods

2.1 Simplified segmentation method with bone probability map

The simplified *bone probability method* is based on the previous method [5], but some functions have been sacrificed to improve processing efficiency. For example, the local phase extraction involving multidimensional inverse fast Fourier transform was removed from *bone probability method* and saved ~ 190 milliseconds. Additionally, some new functions were added to compensate for the accuracy lost due to the deletion. To maximise the performance efficiency, only six amended steps were applied in the present method. Details of each function are described as below (Fig 1):

Bone shadow Due to the high acoustic impedance difference between bone and soft tissue, most US energy is reflected by the bone surface. The region below the bone surface usually appears to be dark in a US image. In the original algorithm [5], the shadow value of the $pixel(a, b)$ was calculated by a Gaussian weighted accumulation of the pixels below. However, considering the fact that the variance ($\sigma^2 < 1$) has been reduced by using a hardware pre-set, the calculation can be simplified as follows:

$$\exp\left(-\frac{(x-1)^2}{2.\sigma^2}\right) \simeq \begin{cases} 1 \text{ si } x = 1 \\ 0 \text{ sinon.} \end{cases} \tag{1}$$

and the *Bone_shadow* can then be generated:

$$Sum(r, c) = \sum_{i=r}^{R} I(i, c) \tag{2}$$

$$Mask(r, c) = 1 - \frac{Sum(r, c) - min_{r,c}(Sum)}{max_{r,c}(Sum) - min_{r,c}(Sum)} \tag{3}$$

$$Bone\_shadow = Mask^{Power} \tag{4}$$

Where r and c are the rows and column indexes in the US image, R is the number of rows in the US image, and I is the pixel intensity matrix. The parameter $Power$ needs to be adjusted and introduced to manage the $Mask$ before generating Bone shadow. Figure 2b demonstrates a bone shadow mask.

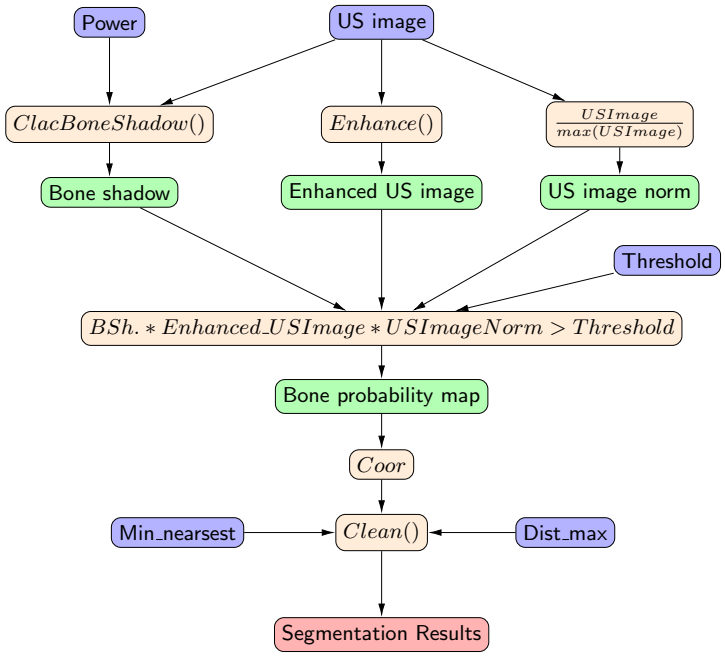

Fig. 1: Flowchart of the method developed in this article

Enhance function An additional function $Enhance()$ has been used to enhance the US image before generating a bone probability map. This function calculated the derivative of the echo intensity, which measured the sensitivity to the echo intensity change. The output image has enhanced edges of the high intensity area, where the soft tissue features were suppressed, and the intensity of bone edges remains at the same level. Figure 2b is an example of the enhanced ultrasound image, and Figure 2e demonstrates a selected column of the echo signal beam ($columnindex = 222$). The peak of the signal ($rowindex = 50$) remains the same level, and the density of reflective artefact ($rowindex = 118$) was reduced. $\forall\ c$, if $r = 1$ or $r = R$,

$$Enhanced\_US\_image(r, c) = 0 \tag{5}$$

$else$,

$$Enhanced\_US\_image(r, c) = 1 - (2 * I_{(r,c)} - I_{(r+1,c)} - I_{(r-1,c)}) \tag{6}$$

Where c is the column index, R is the number of rows in the US image, and I is the pixel intensity matrix.

Bone probability The final bone probability map was generated by multiplying the bone shadow map, the enhanced US image and the normalised US

image. A *threshold* value was applied to eliminate pixels with low probabilities. This parameter is adjustable to improve segmentation quality, and it may need to be adjusted for specific acquisition protocols (depending on the probe and acquisition mode). A higher threshold value can decrease the likelihood of labelling high intensity features, such as tendon, as bony areas; Therefore, with a very high threshold value, the whole bone might remain undetected; and noisy labels will be generated when applied with a very low threshold value. The ideal way of setting the threshold value was to extract the echo intensity value from a "boneless" region.

Segmentation According to the principle of ultrasound beam reflection, only one pixel per column on the bone probability map (column c, row s_c) is recognised to initiate the segmentation (Eq. 7). The recognition was performed as follows:

$$\forall c, \exists s_c \in [0, R], s = \sum_{i=r}^{R} i * Bness(i, c) \tag{7}$$

Where c is the column index, R is the number of rows in the US image, and $Bness$ is the bone probability map described above. After generating the bone probability map, the segmentation can then be obtained as follows:

$$Segmentation'(r, c) = \begin{cases} 1 \text{ if } \quad r = s_c \\ 0 \text{ else} \end{cases} \tag{8}$$

Clean function A $Clean()$ function has been created to improve the segmentation and remove parasite pixels. This step was based on user observation of the preliminary segmentation label (Figure2f). Mislabelled pixels were usually isolated from the main cluster of correct labels. Typically, the pixel density around the false positives area was lower, and it can be examined by using the process below:

$$Sum'(r, c) = \sum_{|j-c|=0}^{2.dist\_max} \left(\sum_{|i-r|=0}^{dist\_max} Segmentation'(i, j) \right) \tag{9}$$

$$Segmentation(r, c) = \begin{cases} 1 \text{ if } \quad Sum'(r, c) > min \quad nearest \\ 0 \text{ else} \end{cases} \tag{10}$$

The new method gave up the principle that was applied in the original method, which is connecting up nearest neighbour segmented pixels into a continuous line, and the $Clean()$ function considered the labels as a cluster of the point cloud. Clusters with continuous shape were included, and the random clusters were excluded. The user can monitor the performance and adjust two parameters to control the level of cleaning. The result of the clean function is demonstrated in Figure 2g.

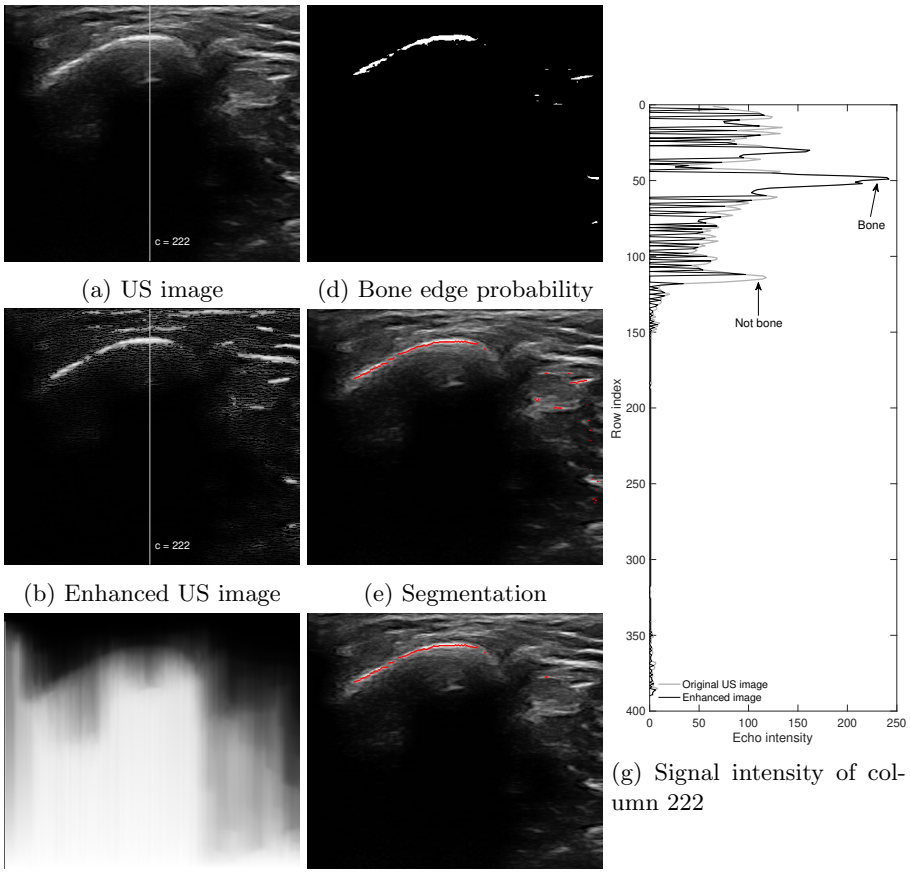

(a) US image (d) Bone edge probability

(b) Enhanced US image (e) Segmentation

(c) Bone shadow (f) cleaned Segmentation

(g) Signal intensity of column 222

Fig. 2: An example of processing results. (a) US image; (b) Enhanced US image; (c) Bone shadow; (d) Bone edge probability map; (e) Segmentation label before clean; (f)Cleaned segmentation label; (g) An example for the *Enhanced*() function, where the y axis is depth of the US image (row index) and the y axis is the intensity of selected column ($c = 222$). the grey line represents the original image and the black line represents the enhanced US image

2.2 Image acquisition and hardware pre-sets

A total of 1393 B-mode US images were collected using a hand-held portable linear US scanner (L7, Clarius Mobile Health CO., Canada) with probe frequency of 4-13 MHz from two young volunteers by sweeping the greater trochanter, the medial femur epicondyle, the medial tibial surface and the radial tuberosity with two hardware pre-sets of the scanner. Each US image was 458×480 pixels, and the depth settings and image resolutions were 4 cm and 0.086 mm. Collected images were independently reviewed, and two experienced practitioners manually segmented the bone surfaces for validation purposes.

To maximise the performance for automatic bone edge recognition, several manufacturer's pre-sets were tested, and one set of parameters originally designed for ocular examination was selected for the further experiments to test the automatic segmentation method. A few mechanical indexes have been adjusted in comparison with the common b-mode settings (Table 1). For example, the derated peak rarefactional pressure associated with the transmit pattern was reduced from 2.66 to 0.487 megapascals; and the pulse repetition frequency was reduced from 9600 to 4800 Hz. The ultrasound power of the selected mode was also reduced from 9.33 to 0.153 milliwatts. Figure 3 demonstrates an example of scanning the same anatomical feature with different hardware parameters. Although scanning muscle and bones with the adjusted mode designated for ocular examinations would suppress soft-tissue resolutions, it saved 50 milliseconds and produced a better accuracy of bone edge labelling in our preliminary testing. Therefore, comprehensive quantitative comparisons between scanning modes were performed in the current study.

2.3 Algorithm testing

The performance of the simplified *bone probability method* was compared against the original *local phase method* [5] and the manual segmentation results. In order to reduce the influence of human error, the average position in each column of the image where both reviewers had labelled a pixel as the bone edge was defined as ground truth data.

Experiments were carried out on a Mac computer with Intel Core i7 CPU, 16 gigabyte RAM, and AMD Radeon Pro 5300M graphic processor and the testing environment was MATLAB (R2019b Update 4). A total of 1393 ultrasound images were independently processed by both the *local phase method* and *bone probability method* segmentation algorithms on the same computer to determine the segmentation speed. The processing time of each image was timed and recorded with *MATLAB function tic, toc*. Two-way ANOVA tests were applied to examine the difference in mean processing times across total images between methods and acquisition modes. The results of both methods were compared with a manually labelled data set using sensitive and specific tests to evaluate the performance. The Accuracy, Recall, Precision, Specificity and F1 score were calculated to exam the performance. Sørensen–Dice coefficient between each automatic and manual label was computed to evaluate the similarity of the shape.

Table 1: Acoustic output parameters comparison between pre-sets

Parameters	Unit	B-mode	Adjusted mode
Mechanical Index			
Pr3[a]	MPa	2.66	0.487
Axial distance	cm	1.90	0.90
Centre frequency	MHz	7.04	6.79
Pulse duration	μsec	0.180	0.183
Pulse repetition frequency	Hz	9600	4800
Pr at PIImax[b]	MPa	4.21	0.601
The soft tissue thermal Index			
Ultrasonic power	mW	9.33	0.153
Active aperture dimensions	cm	1.34	0.499
Centre frequency	MHz	7.04	6.79
Focal length (parallel)	cm	2	0.9
Focal length (vertical)	cm	4	4

[a.] *peak rarefactional pressure (in MPa) that has been derated by 0.3 dB/MHz-cm*

[b.] *Peak rarefaction pressure at maximum pulse intensity integral*

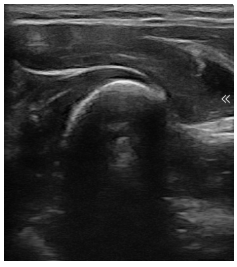

(a) MSK mode image

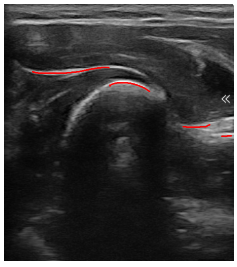

(c) MSK mode with label

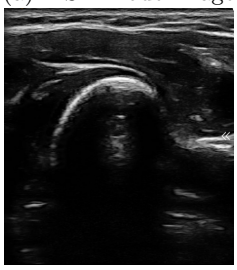

(b) Adjusted mode image

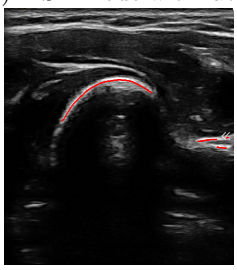

(d) Adjusted mode with label

Fig. 3: An example of comparison between hardware pre-sets. a)MSK mode, b)Adjusted mode, c)labelled MSK image, d)labelled adjusted image (segmentation labels were generated with the local phase method [5])

All performance metrics between methods and acquisition modes were compared with two-way ANOVA.

2.4 Performance testing against a machine learning model

To explore the performance comparison against machine learning frameworks for image segmentation, we compared the bone probability method with a 2D U-Net model [11] trained with a separate set of 235 images. Among these images, 200 was used for training, and 35 was used for validation. Images were normalised and resized to 256×256 pixels. The code was run in Python using TensorFlow with an Ubuntu 20.04 Operating System on an Nvidia GeForce RTX3090 GPU with 24GB memory and Intel i9-10900K CPU. After training, an independent set of 15 images was used for the performance test. The results of *bone probability method* and 2D U-Net were compared with the contour masks used in the training process. The same set of evaluation scores for assessing the quality of the segmentation between methods were computed, and paired t-tests were performed to exam the systematic difference. Segmentation times were not compared because two methods were performed on different machines, and the image size was adjusted to a smaller scale in this sub-experiment.

3 Results

3.1 Processing time

On average, the *bone probability method* took 17.6 ± 5.5 and 19.9 ± 8.1 milliseconds to segment one US image acquired with standard MSK mode and adjusted mode, while the original *local phase method* needed 327.2 ± 14.1 and 335.4 ± 14.2 milliseconds to process the image. The *bone probability method* was significantly quicker than the *local phase method* ($p < 0.01$). Both methods needed less time to process images acquired with MSK mode ($p < 0.01$). The interaction between two factors (processing methods and acquisition methods) was significant ($p < 0.01$)

3.2 Quantitative comparison between methods

The results from both automatic segmentation methods were compared to the manual segmentation. Table 2 shows the results of the comparison with both the *local phase method* and the *bone probability method*. All included images contained bony anatomy; the segmentation results labelled no bone in the image were considered "detection failure". There was no detection failure reported from this particular set of data. As can be seen from Table 1, despite the difference being small, *local phase method* was statistically better than *bone probability method* in Accuracy, Specificity, Precision, Dice and F_1 Score ($p < 0.01$). Notably, *bone probability method* has higher recall metrics when labelling images acquired with adjusted mode ($p < 0.01$).

The comparison between common musculoskeletal mode and adjusted mode can also be seen in Table 2. The ANOVA showed that both methods had better performance in processing images acquired with the adjusted acquisition mode.

Table 2: Classification table comparing *Local Phase method* and *Bone Probability method* against the manual segmentation(with the best performance scores hightlighted in red).

	Local Phase method		Bone Probability method	
AcqMode	MSK	Adjusted	MSK	Adjusted
Accuracy	0.9993(0.0005)	0.9993(0.0005)	0.9984(0.0007)	0.9985(0.0006)
Recall	0.6459(0.2297)	0.6139(0.2550)	0.4520(0.2591)	0.6663(0.2441)
Specificity	0.9998(0.0003)	0.9998(0.0003)	0.9991(0.0005)	0.9989(0.0004)
Precision	0.7885(0.2384)	0.8110(0.2758)	0.5539(0.2937)	0.6688(0.2533)
Dice	0.6675(0.2142)	0.6754(0.2510)	0.4205(0.2181)	0.5722(0.1960)
F_1	0.7025(0.1868)	0.7180(0.2104)	0.5324(0.2171)	0.6580(0.2309)

values: mean (standard deviation)

3.3 Qualitative comparison between methods

The qualitative results of two tested methods are shown in Figure 4, where the red pixels indicate bone labels produced by *local phase method*, while magenta pixels indicate the labels by the *bone probability method* for the segmentation. Bone segmentation results against experienced manual labels are presented side by side in the figure. Two acquisition settings for B-mode US data were also presented for comparison. First and second columns were the same image captured with standard MSK mode, and the third and fourth columns were images captured with the adjusted acquisition mode.

Investigating the segmentation results, we can infer that when the images were captured with the standard MSK settings, the proposed *bone probability method* did not perform as well as the *local phase method*:

- the small gap from the manual labels
- missing bone boundaries
- false bone labels when strong reflection occurred on the sharp bony surface

Although the performance increased when the images were captured with the adjusted mode, some false positive segmentation is still visible. On the other hand, the original method using *local phase feature categorisation* performed better with the adjusted acquisition mode. No false-positive segmentation was observed, despite some false negative being observed. Qualitatively, the *bone probability method* achieves similar performance for the data set that captured with adjusted mode.

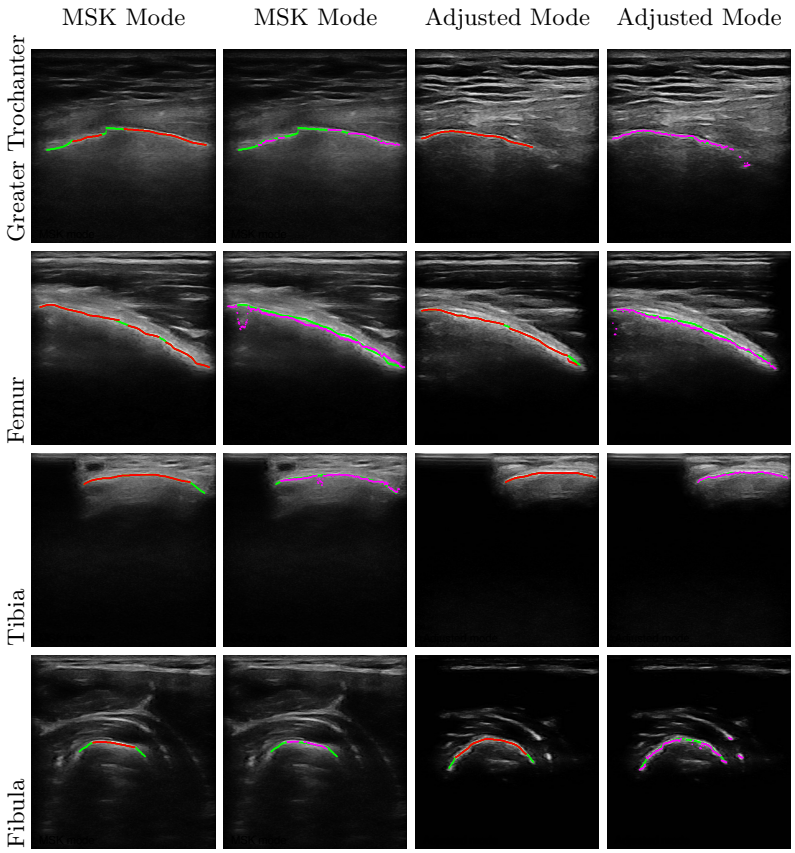

Fig. 4: Bone labels obtained from the *local phase method* (red) and proposed *bone probability method* (magenta) to manual segmentation (green).

3.4 Performance comparison with U-Net

The evaluation results of segmentation methods with and without machine learning were compared to the manual segmentation. Table 3 shows the numerical results of the comparisons. The *Bone probability method* achieved a same level of specificity($p = 0.65$) with the 2D U-NET model, and a better recall metric ($p < 0.01$). However, the performance of proposed method was slightly weaker in Accuracy ($p < 0.01$), Precision ($p < 0.01$), Dice ($p < 0.01$) and F_1 Score ($p = 0.02$), despite the difference being small.

4 Discussion and Conclusion

In this study, we attempted to use a simplified *bone probability method* for bone segmentation from US data. Our previous method [5] that incorporated local

Table 3: Classification table comparing machine learning (2D U-NET) method and *Bone Probability method* against the manual segmentation (with the better score hightlighted in red).

Methods	Accuracy	Recall	Specificity	Precision	Dice	F1
BP method	0.9987	0.9843	0.9987	0.7846	0.7232	0.9161
2D U-NET	0.9995	0.9596	0.9996	0.9643	0.8843	0.9725
p value	< 0.01	< 0.01	0.65	< 0.01	< 0.01	0.02

phase images in conjunction with B-mode US data achieved an acceptable level in quality bone edge localisation. However, the processing time is a significant barrier to the clinical application that we are currently developing. We have investigated how to achieve a similar quality of bone labelling without combining information from local phase images and B-mode US data since it is the most time-consuming step. Our result demonstrates that the proposed *bone probability method* successfully reduced the processing time from ~ 335 ms to ~ 19 ms per image in the size of 458×480 pixels. This is due to the use of bone probability map for effectively preventing the bone shadow region from being processed, while the original *local phase method* had to process the entire image. This processing speed enables us to perform ~ 50 bone edge segmentation per second, which promises real-time data processing while capturing the B-mode US data. This was an essential improvement for the proposed kinematic assessment system with ultrasound.

Although the computational time performance has been significantly improved by giving up the local phase extraction, the qualitative performance assessment was slightly unsatisfied. For our primary purpose of building an automatic 3D bone surface reconstruction method and kinematic assessments, the most crucial performance metrics are specificity and precision. High specificity means fewer false-negative labels that create ghost points when transforming 2D points into the 3D coordinate system and require extra cleaning work. Similarly, high precision means a greater number of true-positive labels, reducing the total time required for scanning and building the completed surface. The original *local phase method* has achieved an adequate level of specificity regardless of the acquisition mode (mean 0.9998, SD:0.0003), and the *bone probability method* was able to achieve a similar level at 0.9989 (SD: 0.0004) with the data set acquired with adjusted mode. However, the *bone probability method* could not improve Precision, although this method did have the best Recall metric across all the conditions. Overall, the *bone probability method* did not perform better than the *local phase method* quantitatively. These outcomes suggest that local phase extraction may still be required to obtain the same level of accuracy, but a faster alternative method is desired.

The high computational time required for the local phase image features extraction is a noticeable drawback for real-time bone edge segmentation. Alsinan et al. [1] reported that it takes 1 second on average to perform local phase

extraction. Our original method was also bottlenecked by the same process, although its execution time can be improved by using more powerful computing hardware. Minimising computational bottleneck can also be achieved by parallelisation. In particular, GPGPUs (general-purpose graphics processing units) are highly suited to the computations commonly found in segmentation algorithms [13]. For operations with low data dependency, such as bone shadow estimation, US image enhancement and normalisation, the throughput may be increased by a factor of 10. This can significantly decrease the time consumed by these bottlenecks.

During this work, the manual segmentation was performed by expert observation and hand labelling. It is worth noting that the precision of manual segmentation is highly dependent on a combination of experience and imaging quality. Some may argue that the huge amount of the laborious task for labelling image contours for machine learning frameworks cannot be a reason to avoid the latest method for image segmentation. However, the most critical barrier for this specific type of task remains to identify ground truth. Even though we can always deploy quantitative metrics to evaluate the segmentation results, it does not entirely represent the quantitative observations. Depending on the application purpose, gold standard annotation is not always present, or not even a very carefully constructed ground truth. A method that can produce a good-enough segmentation, which identifies and validates bone tissue from ultrasound in the absence of reliable ground truth, is still desired.

Although applying a learning free method has a few advantages, the robustness of this type of algorithm remains a challenge. When comparing with the 2D U-net model, the *bone probability method* was able to achieve a similar level of specificity at 0.9987 (2D U-Net: 0.9996) and a better recall score at 0.9843 (2D U-Net: 0.9596). However, the overall scores of the learning free method were not better. Currently, researchers are still attempting different frameworks to achieve a better performance, for example, deploying a two-stage design for U-Net [3], or a novel generative adversarial network architecture [2]. the basic image preset and adjustment mentioned in the present work may also be considered when design a new machine learning framework.

The present study attempted to use the simplest steps to achieve a similar level of segmentation tasks. The results should not be treated as the reason for avoiding further machine learning framework or any advanced image processing method in such tasks. A fast and robust method that can be deployed on a consumable machine is required. Our future work will involve (1) extensive validation of bone labelling against known geometry; (2) improving the accuracy of the *bone probability method* by adding more functions or attempting to reduce the computational cost of local phase feature extraction, (3) replacing the local phase extraction part by using the mask produced by our *bone probability method* and combine with fusion network models proposed by Ronneberger et al. [11] and Alsinan et al. [1].

References

1. Alsinan, A.Z., Patel, V.M., Hacihaliloglu, I.: Automatic segmentation of bone surfaces from ultrasound using a filter-layer-guided cnn. International Journal of Computer Assisted Radiology and Surgery **14**(5), 775–783 (2019). https://doi.org/10.1007/s11548-019-01934-0
2. Alsinan, A.Z., Patel, V.M., Hacihaliloglu, I.: Bone shadow segmentation from ultrasound data for orthopedic surgery using gan. International Journal of Computer Assisted Radiology and Surgery **15**(9), 1477–1485 (2020). https://doi.org/10.1007/s11548-020-02221-z
3. Amiri, M., Brooks, R., Behboodi, B., Rivaz, H.: Two-stage ultrasound image segmentation using u-net and test time augmentation. International Journal of Computer Assisted Radiology and Surgery **15**(6), 981–988 (2020). https://doi.org/10.1007/s11548-020-02158-3
4. Foroughi, P., Boctor, E., Swartz, M.J., Taylor, R.H., Fichtinger, G.: P6d-2 ultrasound bone segmentation using dynamic programming. In: 2007 IEEE Ultrasonics Symposium Proceedings. pp. 2523–2526. IEEE (2007). https://doi.org/10.1109/ULTSYM.2007.635
5. Jia, R., Mellon, S.J., Hansjee, S., Monk, A.P., Murray, D.W., Noble, J.A.: Automatic bone segmentation in ultrasound images using local phase features and dynamic programming. In: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI). pp. 1005–1008 (2016). https://doi.org/10.1109/ISBI.2016.7493435
6. Jia, R., Monk, A.P., Murray, D.W., Mellon, S.J., Noble, J.A.: Greater trochanter tracking in ultrasound imaging during gait. In: 2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI). pp. 260–263 (2015). https://doi.org/10.1109/ISBI.2015.7163863
7. Jia, R., Mellon, S., Monk, P., Murray, D., Noble, A.: Globally optimal registration for describing joint kinematics. Procedia Computer Science **90**, 188 – 193 (2016). https://doi.org/10.1016/j.procs.2016.07.016, 20th Conference on Medical Image Understanding and Analysis (MIUA 2016)
8. Jia, R., Mellon, S., Monk, P., Murray, D., Noble, J.A.: A computer-aided tracking and motion analysis with ultrasound (cat & maus) system for the description of hip joint kinematics. International Journal of Computer Assisted Radiology and Surgery **11**(11), 1965–1977 (2016). https://doi.org/10.1007/s11548-016-1443-y
9. Jia, R., Monk, P., Murray, D., Noble, J.A., Mellon, S.: Cat & maus: A novel system for true dynamic motion measurement of underlying bony structures with compensation for soft tissue movement. Journal of Biomechanics **62**, 156–164 (2017). https://doi.org/10.1016/j.jbiomech.2017.04.015
10. Monk, A.P., Chen, M., Mellon, S., Gibbons, C.L.M.H., Beard, D.J., Gill, H., Murray, D.W.: Measurement of in-vivo patella kinematics using motion analysis and ultrasound (maus). In: 2013 IEEE International Symposium on Medical Measurements and Applications (MeMeA). pp. 257–260 (2013). https://doi.org/10.1109/MeMeA.2013.6549747
11. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015). pp. 234–241. Springer (2015). https://doi.org/10.1007/978-3-319-24574-4_28
12. Salehi, M., Prevost, R., Moctezuma, J.L., Navab, N., Wein, W.: Precise ultrasound bone registration with learning-based segmentation and speed of sound calibration. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI 2017). pp. 682–690. Springer (2017)

13. Smistad, E., Falch, T.L., Bozorgi, M., Elster, A.C., Lindseth, F.: Medical image segmentation on gpus: A comprehensive review. Medical Image Analysis **20**(1), 1 – 18 (2015). https://doi.org/10.1016/j.media.2014.10.012, https://www.sciencedirect.com/science/article/pii/S1361841514001819
14. Villa, M., Dardenne, G., Nasan, M., Letissier, H., Hamitouche, C., Stindel, E.: Fcn-based approach for the automatic segmentation of bone surfaces in ultrasound images. International Journal of Computer Assisted Radiology and Surgery **13**(11), 1707–1716 (2018). https://doi.org/10.1007/s11548-018-1856-x
15. Wang, P., Patel, V.M., Hacihaliloglu, I.: Simultaneous segmentation and classification of bone surfaces from ultrasound using a multi-feature guided cnn. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 134–142. Springer (2018). https://doi.org/https://doi.org/10.1007/978-3-030-00937-3_16