

Spot the Best Frame: Towards Intelligent Automated Selection of the Optimal Frame for Initialisation of Focal Liver Lesion Candidates in Contrast-Enhanced Ultrasound Video Sequences

Spyridon Bakas, Gordon Hunter, Dimitrios Makris
Digital Imaging Research Centre
Kingston University
London, United Kingdom
{S.Bakas, G.Hunter, D.Makris}@kingston.ac.uk

Célia Thiebaud
Télécom Physique Strasbourg
Strasbourg, France
Celia.Thiebaud@gmail.com

Abstract—This paper describes a contribution to a wider project which aims to provide an intelligent automated assistant to radiologists performing the skilled and time-intensive task of detecting and characterising cancerous lesions within a human liver from Contrast-Enhanced Ultrasound (CEUS) video sequences. This particular contribution relates to automatically locating the optimal frame, for initialising a suspected focal liver lesion (FLL), within a CEUS video sequence. Currently, this task is routinely performed manually by radiologists, but is very time-consuming. The proposed approach is to use statistical and image processing techniques to automatically identify the most suitable frame for performing this initialisation, which should save the radiologist significant time and effort, bearing in mind the continuously increasing amount of CEUS data acquired and processed. In the future, this could be coupled with a method for automatically initialising the FLL's area within the area of the ultrasonographic image in this optimal frame and, together with already produced systems for the tracking and characterisation of such lesions, lead to a fully automated system assisting clinicians in the diagnosis of such lesions.

Keywords—Medical Image Analysis; Contrast-Enhanced Ultrasound; CEUS; Focal Liver Lesions; Region Of Interest Initialisation; Computational Intelligence; Human-Computer Interaction

I. INTRODUCTION

Contrast-Enhanced Ultrasound (CEUS) is a widely used method in Radiology and is accepted to be a particularly useful modality for the detection and characterisation of focal liver lesions (FLLs) [4], which are potentially malignant nodules within the liver. In comparison with conventional Ultrasound, CEUS provides a grey-scale enhanced display by maximising the contrast and the spatial resolution between the FLLs and the healthy liver area (parenchyma) through the use of injected contrast-enhanced agents [5]. More specifically, CEUS offers a peripheral enhancement to nodules, based on the fact that the FLLs are solid- or liquid-containing nodules, foreign to the liver's anatomy. CEUS has improved to the point at which it exceeds the sensitivity and specificity of other modalities, such as computed tomography (CT) and magnetic resonance imaging (MRI),






Figure 1. Schematic examples of the four major “signature signals” for FLL characterisation. In the first column of graphs, the red curves depict the temporal profile of the brightness intensity of the FLL, while the black curves show the corresponding profile for the parenchyma. In the second column, the “signature signals” are depicted, which are derived by subtracting the black curve from the red curve. Different signals correspond to different medical conditions.

for diagnosis of lesions within the liver [4].


There is a sequence of tasks that need to be performed for a valid diagnosis to be made. After the acquisition of a CEUS video sequence, the radiologist studies the sequence ‘offline’ and needs to find the frame where the boundaries of the FLL are clearly visible in the image plane, providing a lesion well-distinguished from the surrounding parenchyma area. Additionally, the identified shape of the FLL should be sufficiently representative, in order to initialise it as the region of interest. Subsequently, the shape of the FLL is followed (“tracked”) throughout the video sequence acquired and its dynamic behaviour (the temporal profile of the brightness intensity) is evaluated in comparison to that of the parenchyma. This provides a signal allowing for the characterisation of the FLL. Different signals correspond to different medical conditions, as shown in Fig. 1.



(a)




(b)



(c)

Figure 2. a) Frame 250 depicts the lesion too early to initialise it, b) Frame 336. Radiologist's decision on the choice of the optimal frame for the FLL initialisation (i.e. gold standard), c) Frame 400, the area surrounding the lesion has become too bright.

The initialisation of the area of the lesion in some frame in the CEUS video sequence is an essential task to be performed, in order to allow the tracking of the FLL through the sequence and consequently its characterisation as either benign or malignant, based on its dynamic behaviour [3]. In order to do this, the radiologist must first find the optimal frame in the video sequence for manually initialising the boundaries of the FLL. Conventionally, working backwards and forwards through the video sequence, the radiologist selects the initialisation frame by eye, choosing the frame where the FLL is depicted with its maximum excitation relative to the remainder of the image plane. This frame is



(a)

(b)

Figure 3. The visible information on a frame; in (a) before the definition of our workspace including the textual information provided for the CEUS operator, and in (b) after its definition, having removed the information outside the ultrasonographic image.

expected to be the one with the maximum contrast between the FLL and the parenchyma, which is noted to occur during the enrichment phase. This will provide well-defined FLL boundaries, allowing for an accurate initialisation.

The task of selecting the frame for the FLL initialisation is very time-consuming for the radiologist and depends on the quality of the video sequence, the CEUS operator, as well as the patient. The quality of the video sequence refers to low signal-to-noise ratio, where the noise arises due to the propagation of the ultrasound (US) waves in soft tissues [6]. Additionally, instability of the clinician's hand when holding the transducer affects the localisation of the FLL within the region of the ultrasonographic image. Moreover, the combination of motion effects arising from the patient's heartbeat and breathing, particularly due to physiopathology (e.g. variability of heart rate, irregular breathing patterns) can affect the acquisition of the video, resulting in potential apparent out-of-plane movement of the FLL. During this movement, the clinician normally tries to manually correct the view of the US target, within the obtained image plane (x, y), by changing the elevation axis (z) of the transducer, but thus introducing an issue of the FLL being dispersed in depth. Furthermore, there are not any standardised criteria for the selection of this frame and therefore the task is dependent on human skill & knowledge and prone to human error.

This project focuses on the development, testing, and evaluation of an interface that will assist a clinician to make a reliable diagnosis whilst intelligently reducing the time and effort required by reducing the potential interactions between the radiologist and the acquired data. We have identified a number of tasks in this process suitable for automation, namely identification of the optimal frame for initialisation, performing the initialisation of the FLL region and motion tracking of this region throughout the video sequence, with the intention of acquiring the temporal profile of its brightness intensity. In this paper, we focus on the task of automatically choosing the optimal frame for the initialisation of the FLL region. Our initial work on motion tracking of a manually initialised FLL and its characterisation based on its dynamic behaviour has already




Figure 4. (a) Average intensity of our workspace over time (in frames). The black curve depicts the original values (noisy representation) and the red curve those after a 'moving average' smoothing is performed. Note the particularly noisy variation up to approximately frame 80. (b) Gradient of smoothed average intensity over time. Maximum is around frame 300. (c) Second derivative of the smoothed average intensity curve, over time. The curve crosses the zero line around frame 300. Initial and terminating frames of the CEUS video sequence are 1 and 491, respectively.

been published elsewhere [1]–[3].

II. PROPOSED METHODS

In order to select the optimal frame for performing the initialisation of the FLL region, we need to understand how an experienced radiologist performs the same task. Therefore, we need to find a method of calculating the contrast between the lesion and the parenchyma and subsequently locate where in the video sequence this contrast reaches its maximum. However, there is the intrinsic problem of this whole procedure being carried out before the region of the

FLL has been identified, forcing us to work with the entire area of the ultrasonographic image (Fig. 3.b).

Our approach takes a subset of the video sequence between specified initial and final frames as input and automatically identifies the optimal frame for FLL initialisation as its output. Input of these initial and final frames is essential since the provided CEUS video sequences included frames acquired before the transducer was focused on the region of interest and thus the beginning and end of each sequence includes irrelevant and incomprehensive data for our purposes. Therefore, in order to reduce errors in this automated selection the irrelevant frames are removed.

We propose and compare three methods for identifying this optimal frame. The first is based on the expectation that maximum contrast will be achieved when the rate of change of the brightness intensity of the image is maximal. The other two investigate how the variation of brightness intensity across the image changes over time - one studying the variation over all pixels in the area of the ultrasonographic image, while the other divides this area into local neighbourhoods of $n \times n$ pixels, where n is small compared with the overall size of the image.

In all methods, the workspace is initially defined automatically as the conical area viewed by CEUS, as described in [3], in order to remove irrelevant information - notably textual data provided for the CEUS operator (Fig. 3). This workspace is then applied as a mask to every frame of the video sequence, selecting only the relevant conical area whilst removing the aforementioned artefacts.

A. Method 1 - Gradient based

In this approach, the average brightness intensity of the workspace is computed for each frame of the video sequence. However, there are very short-term fluctuations in this intensity, some of which are due to noise and/or physiological motion of the patient's inner organs. These fluctuations are irrelevant to the dynamic behaviour of the defined workspace over a longer timescale and hence these average brightness intensity values are smoothed over a short timescale using a moving average filter. This filter smooths data by replacing each data point with the average of it and the neighbouring data points defined over a specified time-window. This is equivalent to a low-pass filter with response given by the formula:

$$y_s(i) = \frac{1}{2N+1} [y(i+N) + y(i+N-1) + \dots + y(i-N)] \quad (1)$$

where, $y(i)$ is the i^{th} data point, $y_s(i)$ is the corresponding smoothed value, N is the number of neighbouring data points to be included on either side of $y_s(i)$ and $2N+1$ is the length of the specified time-window.

An example of this filter in use is shown in Fig. 4.a, where it can be observed that the smoothed average brightness intensity of the workspace is continuously increasing after




Figure 5. The optimal frame for initialisation of the FLL as selected by the radiologist in (a) and as selected by our method in (b). No significant differences can be observed visually that can affect the initialisation procedure.

the 80th frame of the video sequence, in contrast to the unsmoothed curve, which shows many short-term fluctuations.

As noted previously, it is expected that maximum contrast in the ultrasonographic image will occur when the rate of change of average brightness intensity reaches its maximum value. Thus, we need to calculate the first derivative of the smoothed average intensity curve in order to obtain its rate of change (Fig. 4.b). Consequently, the point where this gradient begins to decline, corresponding to the maximal rate of change, can be obtained by finding where the second derivative of the smoothed brightness intensity is zero.

The graph in Fig. 4.b shows that, in this example, the gradient of the smoothed brightness intensity curve shows a maximum, i.e. stops increasing, at approximately frame 300. In order to locate this point more precisely, the second derivative has to be computed (i.e. the “gradient of the gradient”) (Fig. 4.c). Note that both the first and second derivatives are smoothed using the same moving average approach as applied to the brightness intensity curve. At the first point where this second derivative, excluding the first 80 frames as noted above, becomes negative (again at approximately frame 300), the rate of increase of the smoothed average brightness intensity begins to decline. However, according to the radiologist consulted, the optimal frame for initialising the lesion in this video sequence

currently being studied would be frame 336. It should be mentioned that the video sequence has been acquired with a frame rate of 25 frames per second and also that the optimal frame given by the radiologist is obtained by visual inspection. This discrepancy corresponds to approximately 1.4 seconds of real time footage.

In order to locate this point of maximum gradient precisely we have to compare the second derivative values at time t and $t + 1$. However, in order to allow for rounding errors and any possibility of residual noise in the smoothed signals, we do not just require a simple decrease between these gradient values, but instead impose a thresholded decrease. Therefore, we accept that the smoothed gradient has genuinely started to decrease when:

$$g_{t+1} < \alpha \cdot g_t \quad (2)$$

where g_t and g_{t+1} are the gradients at time t and $t + 1$, respectively and α is a thresholding factor, less than but close to 1. The optimal value of α can be determined empirically. The first frame t for which (2) is satisfied, excluding any artefacts at the very start of the video sequence, is considered to be the frame where the maximum contrast between the FLL and the parenchyma occurs.

Specifically for the currently examined example video sequence, (2) is satisfied at frame $t = 340$, whereas the radiologist’s choice of optimal frame was $t = 336$, giving a difference of 4 frames, corresponding to 160 milliseconds. This time difference is less than the variation between choices of different but experienced CEUS operators and is therefore considered insignificant. As shown in Fig. 5, there are no obvious visible differences between the frame selected by the radiologist and the frame automatically selected by our method.

B. Method 2 - Spread of pixel intensities

In this approach, we attempt to find the frame of maximum contrast directly from the “spread” of the brightness intensities of all the pixels (fine-grained resolution). The complete range (i.e. the difference between the maximum and the minimum values) of the brightness intensities over a given frame will tend to be exaggerated due to the influence of a small number of excessively bright and dark pixels, which are irrelevant to the initialisation. In order to avoid this artefact, we investigate using the standard deviation (SD) - see (3) - and inter-quartile range (IQR) of the pixel intensities within the ultrasonographic image of each frame as methods of computing a useful quantitative measure of “contrast”. We propose to locate the frame of best contrast, and hence optimal for FLL initialisation, by studying how these quantities change over time.

$$SD = \sqrt{\frac{1}{W \cdot H} \sum_{\substack{1 \leq i \leq W \\ 1 \leq j \leq H}} (p_{i,j} - \bar{p})^2} \quad (3)$$

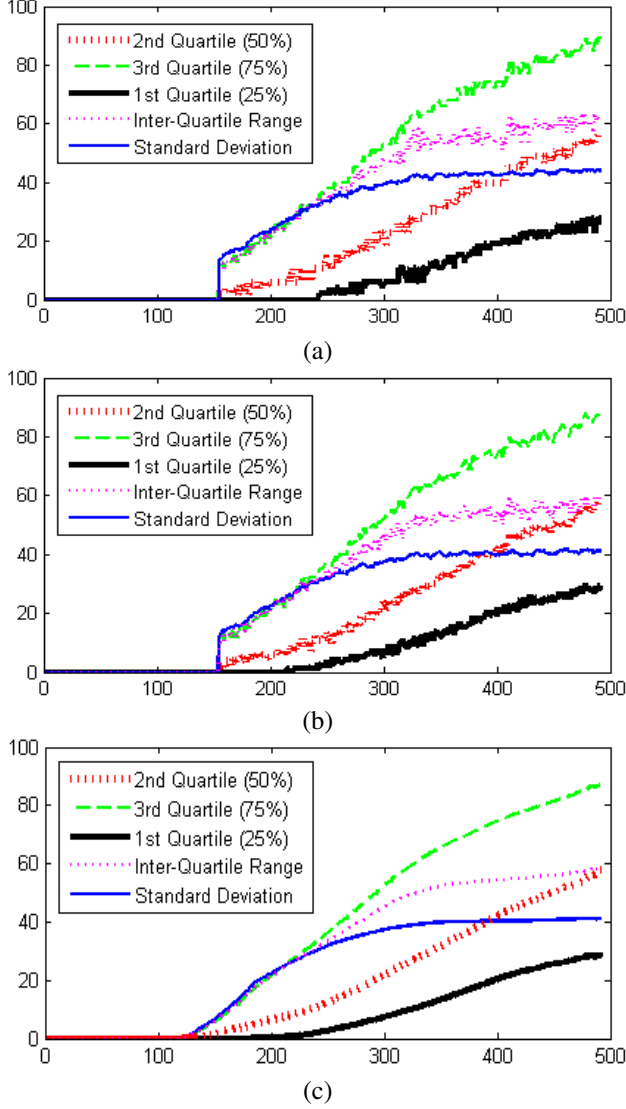


Figure 6. Dynamic behaviour of brightness intensity values. (a) depicts data obtained from Method 2 and the noisy fluctuations are clearly depicted, whereas (b) depicts the data as obtained from Method 3 and (c) depicts the same results as (b) after a moving average filter has been applied to them. The noisy fluctuations are less extreme in graph (b) than in graph (a), due to the averaging over block of pixels (spatial averaging) but still require smoothing over time.

where W and H are the width and the height of the workspace in pixels, respectively, $p_{i,j}$ is the brightness intensity of the pixel at location (i, j) in the current frame and \bar{p} is the mean brightness intensity of this frame, defined as:

$$\bar{p} = \frac{1}{W \cdot H} \sum_{\substack{1 \leq i \leq W \\ 1 \leq j \leq H}} p_{i,j} \quad (4)$$

The IQR is defined to be the difference in brightness intensity values between the first and third quartiles, where these are defined to be the points in an ordered list of pixel brightness intensities such that 25% of all the pixels within

the workspace are darker than the first quartile and 25% are brighter than the third quartile.

An example of this is shown in Fig. 8. If the graphs of SD and/or IQR showed a clear maximum at a particular frame, then this was selected as the optimal one for initialising the FLL. Otherwise, if these quantities continued to increase throughout the video sequence, the point where the SD (or equivalently IQR) reached a fraction β of its final value was chosen as the best candidate frame for the initialisation.

$$SD_{bestFrame} = \beta \cdot SD_{final} \quad (5)$$

where β is close to, but less than, 1. The best value of β was found empirically. In none of the cases studied did the SD and/or the IQR decrease through the early part of the sequence (i.e. they were either continuously increasing or showed a clear maximum).

C. Method 3 - Spread of neighbourhood average intensities

An identified limitation of the aforementioned method (*Method 2 - Spread of pixel intensities*) is that it does not take any account of the location of pixels. In order to find the best frame for initialising the FLL, what is really desirable is to obtain high contrast between different regions of the image - in particular to distinguish between the lesion and the directly surrounding tissue (the parenchyma). In an attempt to achieve this, the workspace is subdivided into local neighbourhoods (“blocks”) of $n \times n$ pixels, where n is small compared with the overall size of the image, and the pixel intensities averaged over that neighbourhood. This provides a coarse-grained resolution of the workspace. Similarly to the aforementioned method (*Method 2*), the SD and IQR of all the neighbourhood average brightness intensities over the workspace are computed for each frame as measures of the contrast between different regions, and how these quantities change over time is investigated. This was carried out for various neighbourhood sizes (i.e. different values of n).

III. EXPERIMENTS

A. Data

The proposed methods are applied and evaluated on real clinical data from seventeen case studies of patients in similar physical condition, provided by the Radiology & Imaging Research Centre of the University of Athens, in Greece. Each case study includes at least one CEUS video sequence (of between 84 and 491 frames) with resolution 768×576 pixels and no compression applied. This dataset included examples from the four major FLL “signals” (i.e. unipolar hyper-enhanced, bipolar hyper-enhanced, unipolar hypo-enhanced, bipolar hypo-enhanced) as shown in Fig. 1.

B. Equipment used

A Siemens ACUSON Sequoia C512 system was used for the data acquisition, equipped with low-frequency 6C2 convex Transducer (2-6 MHz) capturing at a frame rate of

25 frames per second. The contrast medium used was the second generation contrast medium, sulphur hexafluoride microbubbles (SonoVue, Bracco Diagnostics, SpA Milan), which allows excellent depiction of the FLL vascularity and perfusion [7]. Specific acquisition parameters for each case are unknown, as they were set by the radiologist separately for each patient.

C. Methods of Evaluation and Choice of Parameters

To evaluate the results, the mean absolute error (MAE) across the seventeen different cases is calculated for each method, as:

$$MAE = \frac{1}{K} \sum_{1 \leq k \leq K} |x_k - x_k^{(GS)}| \quad (6)$$

where K is the number of cases, x_k is the value obtained for case k by the proposed methods and $x_k^{(GS)}$ is the radiologist's "gold standard" value for case k . The "gold standard" refers to the radiologist's decision on the choice of the optimal frame for the initialisation of the FLL region, which has been based on visual inspection of the CEUS video sequence for each case.

Through experimentation, the optimal choice of the thresholding factor α in (2) was found to be 0.98. Similarly, the optimal value of the thresholding fraction, β in (5), for the SD was found by experiment to be 0.95. Results for this, and also based on the IQR using both $\beta = 0.90$ and $\beta = 0.95$, are shown in Tables II, III and Fig. 8. However, from these tables it is clear that the results based on the IQR were generally inferior to those obtained using the SD. Table I shows the optimal frame obtained for each case using each of the three methods, where the SD approach was employed for each of methods 2 and 3.

Neighbourhood sizes of $n = 4, 8$, and 16 were used in "Method 3". The results obtained for $n = 4$ were almost identical to those from "Method 2". This was as expected since averaging over such a small neighbourhood (just 16 pixels) would not result in much spatial differentiation. Furthermore, although values of $n = 8$ (neighbourhood of 64 pixels) and $n = 16$ (neighbourhood of 256 pixels) yielded similar results in terms of the accuracy of the optimal frame selected, we choose $n = 16$ as the optimal neighbourhood size since this proved to be more time-efficient than $n = 8$.

D. Results

The results obtained, after the application of the three proposed methods to the provided dataset, are compared with the radiologist's "gold standard" in Table I. In Tables II, III and Fig. 8 we show the summary statistics for the errors (relative to this "gold standard") for each method, including the variants of methods 2 and 3 which employed the IQR rather than the SD. The differences between the radiologist's "gold standard" and the automated decisions of the proposed methods might seem to be rather large

Table I
THE OPTIMAL FRAMES. COMPARISON BETWEEN THE AUTOMATED DECISION OF THE PROPOSED METHODS AND THE RADIOLOGIST'S GOLD STANDARD. THE ERROR IN MILLISECONDS IS ALSO PROVIDED.

	Gold Standard	Method 1		Method 2		Method 3	
Case	Frame	Frame	Error(ms)	Frame	Error(ms)	Frame	Error(ms)
1	336	340	-160	344	-320	330	+240
2	260	290	-1200	428	-6720	284	-960
3	315	280	+1400	390	-3000	387	-2880
4	247	257	-400	372	-5000	158	+3560
5	162	180	-720	175	-520	176	-560
6	276	237	+1560	523	-9880	290	-560
7	98	120	-880	197	-3960	129	-1240
8	87	83	+160	110	-920	110	-920
9	139	114	+1000	127	+480	126	+520
10	156	202	-1840	294	-5520	161	-200
11	343	342	+40	423	-3200	422	-3160
12	471	476	-200	476	-200	476	-200
13	216	188	+1120	281	-2600	280	-2560
14	332	65	+10680	218	+4560	214	+4720
15	130	114	+640	144	-560	155	-1000
16	399	354	+1800	434	-1400	375	+960
17	395	290	+4200	413	-720	371	+960

Positive values in *Error (ms)*, indicate that the methods gave an earlier frame, and negative values indicate that the methods gave a later frame, in the CEUS video sequence, than specified by the "gold standard". The distribution of positive and negative errors for each method indicates that none shows a systematic error.

Table II
MEAN AND MEDIAN ERRORS IN FRAMES AND MILLISECONDS FOR EACH METHOD.

	MEAN ERROR		MEDIAN ERROR	
	(frames)	(ms)	(frames)	(ms)
Method 1	+25.29	+1012	+4	+160
Method 2 (SD 95%)	-56.06	-2242	-35	-1400
Method 2 (IQR 90%)	+73.29	+2931	+58	+2320
Method 2 (IQR 95%)	+50.59	+2023	+65	+2600
Method 3 (SD 95%)	-4.82	-192	-14	-560
Method 3 (IQR 90%)	+14.75	+590	+36	+1440
Method 3 (IQR 95%)	-36	-1440	-48	-1920

Positive values indicate that the methods gave an earlier frame, and negative values indicate that the methods gave a later frame in the CEUS video sequence, than specified by the "gold standard".

Table III
SUMMARY RESULTS: MEAN ABSOLUTE ERROR (MAE), IN FRAMES AND MILLISECONDS FOR EACH METHOD.

	MEAN ABSOLUTE ERROR	
	(frames)	(ms)
Method 1	41.18	1647
Method 2 (SD 95%)	72.88	2915
Method 2 (IQR 90%)	87.76	3510
Method 2 (IQR 95%)	71.88	2875
Method 3 (SD 95%)	37.06	1482
Method 3 (IQR 90%)	58.35	2334
Method 3 (IQR 95%)	82.88	3315

The MAE values for each method are graphically depicted in Fig. 8.









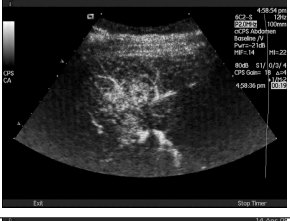
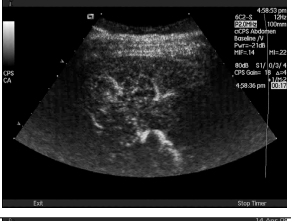

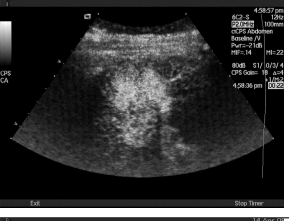












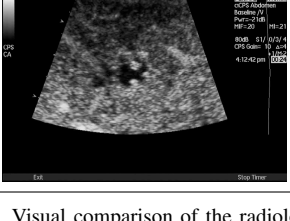
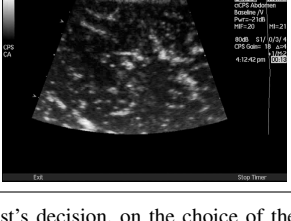


CASE	GOLD STANDARD	METHOD 1	METHOD 2	METHOD 3
1				
2				
3				
4				
6				
10				
14				

Figure 7. Visual comparison of the radiologist’s decision, on the choice of the optimal frame for the initialisation of the FLL region, with the outcome of the proposed methods. The figures shown are the actual video frames from the CEUS video sequences of seven of our case studies. The first column depicts the index of the cases shown, corresponding to the values in Table I. The second column depicts the “gold standard”, whilst the rest (i.e. columns 3-5) depict the decision of the proposed methods.




Figure 8. Summary Results: Mean Absolute Error (MAE) for each method.

- typically up to around 75 frames (3 seconds of the video's acquisition). However it should be noted that there is substantial uncertainty in the precise choice of optimal frame by visual inspection, even by an experienced CEUS operator. Fig. 7 shows the actual video frames, as selected by each of our methods and the "gold standard" decision of the radiologist's choice on the optimal frame for initialising the region of the FLL, from the CEUS video sequences for seven of our case studies, providing a visual comparison of the outcomes. Our results are currently being assessed by the radiologist who provided the original data.

The particularly poor results for case 14 are quite notable in Table I, as well as in Fig. 7. In this example (case 14), the contrast over the whole workspace - and even between the FLL and the parenchyma - shows an early peak around frames 150-160, then remains relatively constant over quite a large part of the video sequence. However, there are dramatic changes occurring in the sequence around frames 240-270 (i.e. dispersion in depth), resulting in a temporary reduction in contrast. The contrast subsequently improves, depicting a secondary peak around frame 340, which is much closer to the radiologist's choice.

IV. CONCLUSIONS AND FUTURE WORK

Based on the aforementioned results, the proposed methods 1 and 3 (particularly using SD rather than IQR) select the optimal frames for initialising the FLL region in reasonable agreement with the choices of the radiologist, bearing in mind that any human judgement by visual inspection will necessarily involve some uncertainty. We conclude that the process of finding the optimal frame for initialising the FLL region can be automated, saving effort on the part of the radiologist and potentially contributing to improving the accuracy and time-efficiency of the whole diagnostic process. However, we observed that the early part of each of the sequences used in our methods contains artefacts not useful to the diagnosis and hence it is essential that the radiologist should specify an initial and a final frame for

each sequence before it is processed.

As noted previously, the work described in this paper is part of a wider project, which aims to fully automate the analysis of CEUS video sequences and assist the radiologist in the diagnosis procedure. The next stages of our work will relate to automatically initialising the region of the FLL and from their temporal profile, subsequently obtained, and further study of their enrichment patterns [7], classify these lesions into the appropriate category, as shown in Fig. 1.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Katerina Chatzimichail and the Radiology & Imaging Research Centre, University of Athens, Evgenidion Hospital, in Greece for providing their real clinical case studies used for the work described in this paper. Spyridon Bakas also thanks the Faculty of Science, Engineering and Computing of Kingston University in London for awarding him a Ph.D. studentship allowing for this research to be conducted.

REFERENCES

- [1] S. Bakas, K. Chatzimichail, A. Autret, A. Hoppe, V. Galariotis and D. Makris, *Localisation and Characterisation of Focal Liver Lesions Using Contrast-Enhanced Ultrasonographic Visual Cues*, In Proceedings of Medical Image Understanding and Analysis (MIUA), King's College, London, UK, 2011
- [2] S. Bakas, K. Chatzimichail, A. Hoppe, V. Galariotis, G. Hunter and D. Makris, *Histogram-based Motion Segmentation and Characterisation of Focal Liver Lesions in CEUS*, Annals of the BMVA, Vol. 2012, No. 7, pp 1-14, 2012
- [3] S. Bakas, A. Hoppe, K. Chatzimichail, V. Galariotis, G. Hunter and D. Makris, *Focal Liver Lesion Tracking in CEUS for Characterisation Based on Dynamic Behaviour*, Springer, Advances in Visual Computing, Lecture Notes in Computer Science, Vol. 7431, pp 32-41, 2012
- [4] M. Claudon, C. F. Dietrich, B. I. Choi, D. O. Cosgrove, M. Kudo, C. P. Nolsoe, F. Piscaglia, S. R. Wilson, R. G. Barr, M. C. Chammas, N. G. Chaubal, M.-H. Chen, D. A. Clevert, J. M. Correias, H. Ding, F. Forsberg, J. B. Fowlkes, R. N. Gibson, B. B. Goldberg, N. Lassau, E. L. S. Leen, R. F. Mattrey, F. Moriyasu, L. Solbiati, H.-P. Weskott and H.-X. Xu, *Guidelines and Good Clinical Practice Recommendations for Contrast Enhanced Ultrasound (CEUS) in the Liver - Update 2012: A WFUMB-EFSUMB Initiative in Cooperation with Representatives of AFSUMB, AIUM, ASUM, FLAUS and ICUS*, Ultrasound in Medicine & Biology, Volume 39, Issue 2, Pages 187-210, 2013
- [5] C. J. Harvey, M. J. K. Blomley, R. J. Eckersley and D. O. Cosgrove, *Developments in Ultrasound Contrast Media*, Eur Radiol, II:675-689, 2001
- [6] J. A. Noble, *Ultrasound image segmentation and tissue characterisation*, Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine 224, pp.307-316, 2010
- [7] S. R. Wilson and P. N. Burns, *Microbubble-Enhanced US in Body Imaging: What role?*, Radiology 257, 24-39, 2010