

Automatic Identification and Localisation of Potential Malignancies in Contrast-Enhanced Ultrasound Liver Scans Using Spatio-Temporal Features

Spyridon Bakas¹(✉), Dimitrios Makris¹(✉), Paul S. Sidhu²,
and Katerina Chatzimichail³

- ¹ Digital Imaging Research Centre, Faculty of Science, Engineering and Computing,
Kingston University, London, UK
{S.Bakas,D.Makris}@kingston.ac.uk
- ² Department of Radiology, King's College Hospital, Denmark Hill, London, UK
PaulSidhu@nhs.net
- ³ Evgenidion Hospital, National and Kapodistrian University, Athens, Greece
Katerina@hcs1.com

Abstract. The identification and localisation of a focal liver lesion (FLL) in Contrast-Enhanced Ultrasound (CEUS) video sequences is crucial for liver cancer diagnosis, treatment planning and follow-up management. Currently, localisation and classification of FLLs between benign and malignant cases in CEUS are routinely performed manually by radiologists, in order to proceed with making a diagnosis, leading to subjective results, prone to misinterpretation and human error. This paper describes a methodology to assist clinicians who regularly perform these tasks, by discharging benign FLL cases and localise potential malignancies in a fully automatic manner by exploiting the perfusion dynamics of a CEUS video. The proposed framework uses local variations of intensity to distinguish between hyper- and hypo-enhancing regions and then analyse their spatial configuration to identify potentially malignant cases. Automatic localisation of the potential malignancy on the image plane is then addressed by clustering, using Expectation-Maximisation for Gaussian Mixture Models. A novel feature that combines description of local dynamic behaviour with spatial proximity is used in this process. Quantitative evaluation, on real clinical data from a retrospective multi-centre study, demonstrates the value of the proposed method.

Keywords: Localisation · Malignancy identification · Contrast-enhanced ultrasound · Focal liver lesion · Liver · Perfusion · Clustering

1 Introduction

Contrast-Enhanced Ultrasound (CEUS) is a modality widely accepted for the detection and characterisation of focal liver lesions (FLLs) [1]. In comparison with

conventional B-mode ultrasound (US), CEUS enhances the contrast between the liver and the FLL through the use of intravenously injected contrast agents [2]. CEUS has diagnostic accuracy higher than 95 % for the evaluation of malignant FLLs [3] and its sensitivity and specificity exceeds that of other modalities, such as contrast-enhanced computed tomography (CE-CT) and contrast-enhanced magnetic resonance imaging (CE-MRI) [4]. It is recognised as the most cost-efficient imaging solution for distinguishing between benign and malignant FLLs [1], after an inconclusive US scan. Also it is in the forefront of CEUS scope to help in reducing the radiation burden to population.

A CEUS liver scan is divided into three different phases over time, namely arterial, portal venous and late phase. These allow the observation of the flow of the intravenously injected contrast agent, by intensity changes in the captured plane. Recording these intensity changes for different tissues during a CEUS sequence leads to time-intensity curves (TICs). These curves describe the perfusion dynamics for different regions and lead to parameters, that allow for the differentiation of the nature of tissues [1]. Specifically, a tissue that shows increased perfusion (hyper-enhancing) in comparison with the rest healthy tissue (parenchyma), during the arterial phase, reveals the typical behaviour of a potentially malignant FLL. This FLL can only be confirmed as a malignancy if during the late phase its region is darker than the parenchyma [5].

Currently, radiologists routinely detect, localise and classify FLLs in CEUS data manually, through a time-consuming series of tasks, namely (i) identification of a reference frame for (ii) localising a region of interest (ROI), e.g. FLL, (iii) monitoring of the dynamic behaviour (i.e. brightness intensity changes) of different ROIs over time, and eventually classifying an FLL as benign or malignant. Each of these tasks requires a high-level of expertise, provides subjective results and is prone to misinterpretation and human error [1]. Processing of CEUS data poses a very challenging task due to intensity changes, acoustic shadows, low signal-to-noise ratio, the transducer movement, as well as the patient's irregular breathing patterns and the motion of the inner human organs that affect the ROIs' apparent 2D size and shapes.

While computer-aided solutions have been suggested for the aforementioned tasks of (i) [6] and (iii) [7–11] in CEUS data, no solution has ever been proposed for the automated localisation of the position and shape of an FLL in a CEUS recording (task (ii)). Considering the continuously increasing number of CEUS scans, automation of the localisation procedure is of much interest as it will assist the diagnostic procedure and the assessment of the repeatability and reproducibility of the examination on different patients or between different clinical centres, as it is expected to provide objective results.

Methods suggested thus far for the task of monitoring FLLs in CEUS (task iii) require the existence of prior manually initialised ROIs (e.g. FLL). A method is described in [10], which first compensates for rigid motion in CEUS video recordings and then displays TICs to assist radiologists to classify an FLL. Motion compensation is based on iterative maximisation of mutual information [12], which is dependent on an intensity constancy constraint. Another method described in [7],

creates a non-rigid motion model by combining histogram information from a ROI with the detection of SIFT keypoints [13]. It then suggests an FLL classification decision based on obtained TICs. Both methods require the manual delineation of the ROI and the part of the image that needs re-alignment (e.g. the conical area of the ultrasonographic image, also named US mask) on a manually identified frame. Then another method, suggested in [8], proposes an affine transformation motion model, based entirely on SIFT keypoints [13], to monitor and classify an FLL based on its TIC. In addition, [8] introduces an automatic delimitation of the US mask based on an intensity change detection. A pixelwise analysis within a manually initialised FLL is performed in [11]. Specifically, after compensating for rigid motion similar to [10], micro-vasculature differences between benign and malignant FLLs are compared and the FLLs' spatial heterogeneity is quantified.

Furthermore, segmentation methods for the localisation of lesions have been applied in CT [14] and MRI [15, 16] data. Specifically, these methods require some manually annotated areas in one slice by an operator, and then they either encode the pairwise similarity of intensity information between different pixels, or model the intensity relationship of each pixel to a set of global intensity clusters, across the 3rd dimension of the data (i.e. different depth slice). Even though the 3rd dimension in volumetric CT and MRI data can be considered relative to the temporal dimension of CEUS data, in CT and MRI the same tissues are represented with consistent brightness intensity across the 3rd dimension, making the localisation task much easier, especially if prior knowledge is included.

The aim of the present study is firstly to identify and categorise hyper- and hypo-enhancing FLLs in the first phase of a CEUS examination, and then to localise regions of potential malignancy within the liver region in CEUS video sequences, in a fully automatic manner and without considering any prior knowledge. This can be ultimately used as an objective initialisation of segmentation methods, as well as to attract the attention of radiologists in the suggested ROIs within the US mask that they might have missed.

2 Method

Specifically, this paper considers the problem of identifying and localising a potential malignancy by clustering together regions with the same dynamic behaviour (TIC) within the US mask, in a similar manner to the process performed conceptually by radiologists. After compensating for the motion included within the sequence of a patient's scan, the US mask is divided into smaller regions and their local TICs are obtained. Subsequently, dimensionality reduction is applied on the TICs for each individual patient, and the principal axes describing more than 90 % of the total variance are used in combination with location information to form the feature vector for each region. A framework based on the optimisation technique of Expectation-Maximisation (EM) with a set of Gaussian Mixture Models (GMM) is then used to segment the US mask into meaningful ROIs based on these feature vectors, and therefore localise the hyper-enhancing regions, i.e. potential malignancies. Figure 1 summarises the whole pipeline of the proposed method.




Fig. 1. Visual summary of the proposed method.

Initially, the US mask is automatically segmented on each patient’s case by considering an intensity change detection, as described in [8], in order to remove all the irrelevant information - notably textual data provided for the CEUS operator. This US mask is then applied to every frame of the sequence, selecting only the relevant information, whilst removing the aforementioned artefacts. Then the optimal reference frame at time t_0 is automatically identified, as proposed in [6]. According to radiologists, this frame is expected to be the one with the maximum contrast between FLL and parenchyma.

Any present motion between the liver and the US transducer affects negatively the processing outcome and therefore the frames of each sequence need to be spatially realigned to compensate for any in-plane movements. Assuming there is only motion within –instead of across– the plane, and a simple translation is sufficient to describe the level of relative motion between the liver and the US transducer, the point-based registration technique of Compact And Real-time Descriptors (CARD) keypoints [17] is employed to automatically estimate the motion between all the frames and realign them according to the reference frame. This is done by first registering keypoints every two successive frames of the acquired sequence and then matching correspondent keypoints by using the Nearest Neighbour Distance Ratio in the descriptor space. The translation between every pair of successive frames is then estimated as the average displacement of the matched correspondences between them. The motion compensated video data can then be processed as a 3D volume, where the 3rd dimension depicts information from the temporal domain.

2.1 Local Time-Intensity Curves

Obtaining information from the brightness intensity of a single pixel (fine-grained resolution) is considered susceptible to speckle noise (i.e. very sensitive to “outliers” – excessively bright and dark pixels). Therefore, the spatial averaging through more coarse-grained resolutions was considered essential, by subdividing the US mask into B non-overlapping local neighbourhoods (“patches”) of $n \times n$ pixels and avoid such an effect, where n is small compared with the overall size of the image. The pixel brightness intensity values are then averaged over each of these patches.

$$\bar{p}_b(t) = \frac{1}{n^2} \sum_{i_b=1}^{n^2} p_{i_b}(t). \quad (1)$$

where $p_{i_b}(t)$ is the intensity of the i^{th} pixel within the b^{th} block at time t .

The changes of the average brightness intensities of each patch (\bar{p}_b) over time (i.e. its dynamic behaviour) can lead to information similar to that used by radiologists. Such information are encompassed in the TICs of the local patches, by estimating the \bar{p}_b for each frame in the sequence.

$$\mathbf{TIC}_b = [\bar{p}_b(1), \bar{p}_b(2), \dots, \bar{p}_b(T)]. \quad (2)$$

where T refers to the number of frames of the CEUS video sequence.

2.2 Early Exclusion of Hypo-enhancement

A foreground mask is produced including only patches where potential malignancies may exist. Specifically, a patch is included in the mask if it is enhancing more than a certain proportion (ψ) of the maximum intensity value of the whole 3D volume. In addition, a morphological opening is used as noise removal, to exclude any patches that are not connected with any other neighbouring patches, or even if connected, the area that they cover is smaller than the size (ϕ) of the structuring element used for the morphological opening.

The resulting foreground mask for a hyper-enhancing case is expected to be solid, as the FLL would have enriched as much as the parenchyma until the end of the arterial phase. On the other hand, for a hypo-enhancing case, this mask is expected to have a ‘hole’ either in the middle or at the side of the mask, depending on where the FLL is located in relation to the parenchyma. This property is exploited to automatically discharge hypo-enhancing cases by separating them from hyper-enhancing cases after applying a specified threshold on the overlap metric of the Jaccard index between the foreground mask and its convex hull.

2.3 Feature Extraction and Clustering

Information of the dynamic behaviour of the foreground patches is essential to be considered in the clustering of the data. However, because T is large (see Sect. 3.1), the data included in each TIC is represented as a high-dimensional vector. Thus, principal component analysis (PCA) by eigenvalue decomposition is employed to reduce the dimensionality of the data and the computational complexity, as well as to create more meaningful feature vectors. The number of dimensions (d) is chosen such that the variance in the data accounted for by these d dimensions (v_d), exceeds a specified proportion, typically 90 %, of the total variance v_T ($v_d > 0.9 \cdot v_T$, for $1 \leq d \leq T$). The coordinates of the centre of each patch (x_{c_b}, y_{c_b}) are also included in each feature vector along with the PCA dimensions of the local TIC, such that each feature vector may represent both the spatial proximity and the dynamic behaviour of the local patch.

The optimisation technique of Expectation-Maximisation (EM) is used with a set of Gaussian Mixture Models (GMM) to perform the clustering of the feature vectors. This is done in an attempt to cluster together patches with similar dynamic behaviour. To find the optimal number of clusters for each case, an overestimated initial number of clusters is used to initialise a Bayesian Classifier with the functionality of a GMM Probability Distribution Function, based on the Figueiredo-Jain algorithm [18].

2.4 Region Selection

Radiologists extract various parameters from a TIC to describe different aspects of a ROI's perfusion and determine its functional features [1]. However, most of these parameters require a continuous sequence throughout the three phases of the exam. Due to the provided data comprising only a short sequence of the exam, including at most a small part of the portal venous phase, the only parameters possible to extract are: the peak intensity (PI), the time to reach the PI (TPI) and the regional blood flow (RBF), which is the area under the TIC from time zero until TPI. RBF is considered as the most useful parameter for the scope of this work as it implicitly includes information about the PI, the TPI and the slope of the TIC.

Sorting the identified clusters in descending order of their RBF is expected to reveal first the cluster with the most hyper-enhancing behaviour. In case of an actual hyper-enhancing FLL this cluster can provide sufficient information for the location, size, shape and nature of the FLL, whilst if it is an actual hypo-enhancing case, then the proposed method is expected to identify regions of the parenchyma that have a hyper-enhancing behaviour in relation to the rest of the image, if not a second FLL. For example, in Fig. 2, cluster depicted by (a) shows the largest RBF, which is typical of a malignant FLL. Furthermore, as the chosen cluster might include some irrelevant pixels (i.e. artefacts) on the image plane (Fig. 2), a proportion of the cluster's pixel population is used




Fig. 2. Clusters obtained from a clinical case, during arterial phase, sorted in descending order of their regional blood flow. *The first row depicts the clusters visualised in the image space and the second row depicts the corresponding TICs, where the vertical and horizontal axes denote brightness intensity and time, respectively. Each of the curves depicts a different patch and the dark dotted curve in the middle of each of the 8 graphs is the average/centroid for each cluster.*

(e.g. the 1st quartile) to automatically provide the location of the potential malignancy on the image plane. This location can ultimately be used as an initialisation seed-point to an existing segmentation method.

3 Experiments and Results

3.1 Materials

The proposed method is evaluated through being applied to real clinical data of 63 case studies, of patients in different physical condition. These cases comprise 34 hyper- and 29 hypo-enhancing FLLs during the arterial phase of a CEUS scan and in most cases information from the portal venous phase was also included. Each case includes one sequence of duration between 3 and 82s. The imaging plane was chosen such that most of the motion is within –instead of across– the plane, allowing for its compensation.

All data was acquired using Siemens ACUSON Ultrasound (US) systems (Mountainview CA). Specifically, 49 cases were captured at King’s College Hospital in the UK, using an S2000 US system equipped with 4 (or) 6 MHz curvilinear transducer at spatial resolution 1024×768 pixels, and the remaining 14 cases were captured at Evgenidion Hospital in Greece, using a Sequoia C512 US system equipped with 6-2 MHz curvilinear transducer, at spatial resolution 768×576 pixels. In all examinations the second generation contrast medium SonoVue [19] (Bracco S.p.A., Italy) was used in a 2.4 ml bolus intravenous injection (into an arm vein). Specific acquisition parameters of the equipment, such as the transducer’s field of view and gain, for each patient are unknown, as they were set by the radiologist individually at the start of each examination. The acquisition method of this data reflects true clinical practice and leads to increased variability. Examinations were performed by radiologists with 13–16 years of experience using CEUS. All data were obtained without prior knowledge of subsequent processing by a software tool and without any specific instructions being given to the radiologist beforehand, hence reflecting true clinical practice. Appropriate ethics and confidentiality procedures have been followed at all times.

3.2 Results

For evaluating the automatic identification of hyper-enhancing (i.e. potentially malignant) cases, the measures of true positive rate (sensitivity) and false positive rate (1-specificity) are used, in the range $[0,1]$. The positive and the negative samples refer to the hyper- and the hypo-enhancing cases, respectively. From the clinical point of view, the true positive (TP) rate should ideally be kept equal to 1, so no single potentially malignant case is missed. At the same time, the proposed method should discharge as many hypo-enhancing cases as possible, keeping the false positive (FP) rate low.

According to [1] the duration of the arterial phase is at least 10s. This justifies our results, shown in Fig. 3(a), where the FP rate is as high as 0.8 (specificity = 0.2) when sequences with acquisition duration (AD) less than 10s are considered. Then for sequences with AD more than 10s, the FP rate drops to 0.46




Fig. 3. Graph (a) denotes values of FP rate (*y axis*) for the acquisition duration of sequences being more than a certain number of seconds (*x axis*), whilst TP rate is equal to 1. Graph (b) shows the ROC curves for cases with duration ≥ 10 & 14 s.

(specificity = 0.54). The best FP rate is equal to 0.38 (specificity = 0.62) and obtained when the AD is at least 14 s. Furthermore, two ROC curves are shown in Fig. 3(b) to assess the possible change of the FP rate in relation to the TP rate, for both AD above 10 and 14 s.

For evaluating the automatic localisation of the potential malignancies on the image plane, first the boundaries of the FLL have been manually annotated by a radiologist in the reference frame providing its ground truth (GT). The performance of the method is then assessed based on whether the location point provided by the proposed method (described in Sect. 2.4) is within the GT, or not. To measure this performance a correct localisation rate is used, over the number of hyper-enhancing cases with AD above 10 and 14 s, separately.

To provide some comparative results, a baseline approach is considered based on automatically thresholding the reference frame, using the Otsu’s method [20]. This is expected to provide a foreground population of pixels that describe mostly the hyper-enhancing FLL. Then, similarly to the proposed method, a proportion of this population is used to obtain the location of the potential malignancy.

The best correct localisation rate achieved is 80 % and obtained for the 40 % of the chosen cluster’s population (Table 1). On the other hand, the best result for the baseline method is 62.07 % after using 40 % of its mask’s population.

Table 1. Correct localisation rate for a seed-point within the actual FLL.

Method	Population’s proportion	Duration thr ≥ 10 ”	Duration thr ≥ 14 ”
Baseline	(25 %)	31.03 %	25 %
Baseline	(40 %)	62.07 %	55 %
Proposed method	(25 %)	51.72 %	60 %
Proposed method	(40 %)	65.52 %	80 %

4 Conclusions and Future Work

This paper demonstrates that the task of identifying and localising a potentially malignant FLL can be performed in a fully automatic manner. Specifically, the proposed framework firstly distinguishes between hyper- and hypo-enhancing cases by using only a video sequence of the arterial phase as input, and then automatically localises potential malignancies on the image plane. The first step is addressed by assessing local intensity variations and analysing their spatial configuration. Then, for the localisation step, a novel feature vector that encompasses the local dynamic behaviour and combines it with the spatial proximity is used in a clustering approach, using EM-GMM.

Experimental results show that the proposed method appears to perform adequately on identifying and localising FLLs with hyper-enhancing behaviour during the arterial phase. Such lesions are of significant importance to radiologists, as they may account for malignancies. FLLs of typical hypo-enhancing behaviour during the arterial phase are benign and therefore of less importance to radiologists.

Further improvements of the proposed framework should include the reduction of the FP rate, making the identification of potential malignancies more specific. Also coupling the proposed localisation approach with an iterative segmentation method might lead to a fully automated and precise approximation of the FLL boundaries, allowing the radiologists to use an automatic delineation of the FLL for its assessment.

References

1. Claudon, M., Dietrich, C.F., Choi, B.I., Cosgrove, D.O., Kudo, M., Nolsøe, C.P., et al.: Guidelines and good clinical practice recommendations for CEUS in the liver - update 2012: a WFUMB-EFSUMB initiative in cooperation with representatives of AFSUMB, AIUM, ASUM, FLAUS and ICUS. *Ultrasound Med. Biol.* **39**, 187–210 (2013)
2. Harvey, C.J., Blomley, M.J.K., Eckersley, R.J., Cosgrove, D.O.: Developments in ultrasound contrast media. *Eur. Radiol.* **11**, 675–689 (2001)
3. Strobel, D., Seitz, K., Blank, W., Schuler, A., Dietrich, C.F., von Herbay, A., et al.: Tumor-specific vascularization pattern of liver metastasis, hepatocellular carcinoma, hemangioma and focal nodular hyperplasia in the differential diagnosis of 1349 liver lesions in contrast-enhanced ultrasound (CEUS). *Ultraschall Med.* **30**, 376–382 (2009)
4. Westwood, M.E., Joore, M.A., Grutters, J.P.C., Redekop, W.K., Armstrong, N., Lee, K., et al.: Contrast-enhanced ultrasound using sonovue®(sulphur hexafluoride microbubbles) compared with contrast-enhanced computed tomography and contrast-enhanced magnetic resonance imaging for the characterisation of focal liver lesions and detection of liver metastases: a systematic review and cost-effectiveness analysis. *Health Technol. Assess.* **17**, 1–243 (2013)
5. Wilson, S.R., Burns, P.N.: Microbubble-enhanced US in body imaging: what role? *Radiology* **257**, 24–39 (2010)

6. Bakas, S., Hunter, G., Thiebaud, C., Makris, D.: 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. In: 9th International Conference on Intelligent Environments, pp. 196–203. IEEE Press (2013)
7. Bakas, S., Chatzimichail, K., Hoppe, A., Galariotis, V., Hunter, G., Makris, D.: Histogram-based motion segmentation and characterisation of focal liver lesions in CEUS. *Ann. BMVA* **2012**, 1–14 (2012)
8. Bakas, S., Hoppe, A., Chatzimichail, K., Galariotis, V., Hunter, G., Makris, D.: Focal liver lesion tracking in CEUS for characterisation based on dynamic behaviour. In: Bebis, G., et al. (eds.) *ISVC 2012, Part I. LNCS*, vol. 7431, pp. 32–41. Springer, Heidelberg (2012)
9. Bakas, S., Sidhu, P.S., Sellars, M.E., Hunter, G.J.A., Makris, D., Chatzimichail, K.: Non-invasive offline characterisation of contrast-enhanced ultrasound evaluations of focal liver lesions: dynamic assessment using a new tracking method. In: 20th European Congress of Radiology (2014)
10. Rognin, N., Campos, R., Thiran, J.P., Messenger, T., Broillet, A., Frinking, P., et al.: A new approach for automatic motion compensation for improved estimation of perfusion quantification parameters in ultrasound imaging. In: 8th French Conference on Acoustics, pp. 61–65 (2006)
11. Ta, C.N., Kono, Y., Barback, C.V., Mattrey, R.F., Kummel, A.C.: Automating tumor classification with pixel-by-pixel contrast-enhanced ultrasound perfusion kinetics. *J. Vac. Sci. Technol. B Nanotechnol. Microelectron.* **30**, 02C103 (2012)
12. Shekhar, R., Zagrodsky, V.: Mutual information-based rigid and nonrigid registration of ultrasound volumes. *IEEE Trans. Med. Imaging* **21**, 9–22 (2002)
13. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* **60**, 91–110 (2004)
14. Zhou, J., Huang, W., Xiong, W., Chen, W., Venkatesh, S.K., Tian, Q.: Delineation of liver tumors from CT scans using spectral clustering with out-of-sample extension and multi-windowing. In: Yoshida, H., Hawkes, D., Vannier, M.W. (eds.) *Abdominal Imaging 2012. LNCS*, vol. 7601, pp. 246–254. Springer, Heidelberg (2012)
15. Crum, W.R.: Spectral clustering and label fusion for 3D tissue classification: sensitivity and consistency analysis. *Ann. BMVA* **2009**, 1–12 (2009)
16. Song, Z., Tustison, N., Avants, B., Gee, J.: Adaptive graph cuts with tissue priors for brain MRI segmentation. In: 3rd International Symposium on Biomedical Imaging: Nano to Macro, pp. 762–765. IEEE Press (2006)
17. Ambai, M., Yoshida, Y.: CARD: compact and real-time descriptors. In: IEEE International Conference on Computer Vision, pp. 97–104 (2011)
18. Figueiredo, M.A.T., Jain, A.K.: Unsupervised learning on finite mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 381–396 (2002)
19. Schneider, M.: Characteristics of sonovue. *Echocardiography* **16**, 743–746 (1999)
20. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man Cybern.* **9**, 62–66 (1979)

Abdominal Imaging. Computational and Clinical
Applications

6th International Workshop, ABDI 2014, Held in
Conjunction with MICCAI 2014, Cambridge, MA, USA,
September 14, 2014.

Yoshida, H.; Näppi, J.J.; Saini, S. (Eds.)

2014, XIII, 296 p. 139 illus., Softcover

ISBN: 978-3-319-13691-2