

Speckle Spot Detection in Ultrasound Images: Application to Speckle Reduction and Speckle Tracking

Nicolas Widynski*, Thierry Géraud[†] and Damien Garcia*

*University of Montreal Hospital,
Montreal, QC, Canada

[†]EPITA Research and Development Laboratory,
Paris, France

Abstract—This paper investigates the speckle spot detection task in ultrasound images. Speckle spots are described by structural criteria: dimensions, shape, and topology. We propose to represent the image using a morphological inclusion tree, from which speckle spots are detected using their structural appearance. This makes the method independent of contrast, and hence robust to intensity correction. The detection was applied to speckle reduction and speckle tracking, and experiments showed that this approach performs well compared to state-of-the-art methods.

I. INTRODUCTION

Speckle is inherent in ultrasound imaging. Depending on the application, speckle can be viewed as noise or as a feature of the underlying tissue. It is considered as noise when viewed as the result of random fluctuations of the backscattered waves. In practice, speckle alters the image by lowering its contrast resolution, which can be a limiting case for the detection of small/low-contrast structures in clinical examinations [1]. Speckle can also be regarded as a feature since its appearance and local echogenicity depend on the number and the locations of the scatterers within the resolution cell. This makes the speckle signal-dependent. In both cases, the characterization of the speckle is an important step of the targeted application. In this article, we address the problems of speckle reduction and speckle tracking. The former sees the speckle as noise, the latter as a feature. Speckle spot detection is employed as the initial step of these two applications.

We propose to detect the speckle spots using a morphological tree representation, namely the tree of shapes [2]. We call speckle *spots* the granular patterns observed in the ultrasound image, also referred to speckle *cells* in the literature [3]. Speckle spots are first defined by structural properties (Sect. II), and are then detected directly from the tree (Sect. III). The tree represents the image by inclusion of connected components. By computing attributes on the tree, the objective is to detect subtrees that correspond to speckle spots. The tree can be filtered, or *pruned*, in order to achieve speckle reduction (Sect. IV). Also, speckle spots can be extracted and thus serve as markers for speckle tracking using a point set registration algorithm (Sect. V). Finally, the use of morphological trees in ultrasound images offer several perspectives (Sect. VI).

II. CHARACTERIZING THE SPECKLE SPOTS

Speckle results from interferences in the backscattered waves produced by punctual random scatterers present in a resolution cell. The coherent summation of the echoes may be constructive or destructive depending on the organization within the resolution cell, and along with the point spread function low-pass filtering, they produce small bright and dark spots in the ultrasound B-mode envelope. Wagner et al. [3] proposed a theoretical prediction of the speckle size using second-order statistics in the presence of fully developed speckle (Rayleigh case). They define the speckle size as related to the autocovariance functions (ACVF) in the lateral and depth directions. They also showed that the correlation length in the Rayleigh scattering media was comparable to the resolution cell size. Bashford and von Ramm [4] conducted an empirical study on the appearance of bright speckle spots in three dimensions. An analysis of their contour maps revealed a structure of concentric elliptical lines, thus forming local maxima, described by the authors as “slices of an onion, with different layers representing different isosurfaces”. When viewing the image as a topological three-dimensional map, bright and dark speckle spots outline peaks and troughs in the landscape. One can thus express the following hypotheses regarding the structure of the speckle spots:

- H1: Speckle spots are connected components including exactly one local maximum/minimum.*
- H2: The shape of the speckle spots can be approximated by ellipses.*
- H3: Speckle spots have dimensions measurable by the ACVFs functions.*

These hypotheses shall serve as motivations of the proposed detector. Notice that the underlying definition of a speckle spot is purely topological. In particular, no assumption has been made on the intensity values. This trait is particularly relevant on ultrasound images, since the intensity range depends both on the acquisition system and non-linear transformations. It is thus expected that structural information will be well preserved under the various compression transformations. The last hypothesis requires to estimate the speckle size. This is done by taking the full-width at half-maximum (FWHM) of the lateral and depth ACVFs [3], [4]. The ACVFs were computed on a manually defined window located at the focus range and presenting homogeneous speckles.

III. DETECTING THE SPECKLE SPOTS




Fig. 1. Two morphological representations of the same image.

We propose a methodology aiming at detecting connected components respecting the geometric and inclusion properties formulated through *H1-3*. Mathematical morphology can help since it is based on set theory, and can study the objects in function of their shape, size, and grayscale information. The tree of shapes, originally introduced in [2], is a morphological self-dual representation of images. Basically, a subtree represents a connected component, and a parent-child node relationship expresses a morphological inclusion. The self-duality property makes the method independent of contrast, which is particularly relevant in the case of ultrasound images as the contrast is hardware dependent. The tree of shapes can also be seen as a merging of the min-tree and the max-tree. Min-tree and max-tree are dual component trees defining inclusions in terms of grayscale levels (Fig. 1).

A. Computing the Morphological Trees

The morphological trees can be computed using a methodology adapted from the Tarjan’s *union-find* algorithm. In this work, we followed the implementation of the min/max-trees proposed in [5], which results in an efficient algorithm with a compact tree structure. Géraud et al. [6] showed that this implementation can be viewed as a “canvas”, meaning that is generic enough to serve different aims. In particular, it can be used for building the tree of shapes. Briefly, the union-find algorithm partitions a set of elements into disjoint subsets. In our application, the disjoint subsets are hierarchically organized and thus define a tree. Each element is a pixel, and represents a node of the tree. The partition is carried out by performing two successive operations for each neighbor of each element: 1) *find* determines if the element and his neighbor belong to the same subtree; if so, 2) *union* hooks up the neighbor subtree to the element.

The procedure is described in Algorithm 1. The auxiliary structure *zpar* stores temporary root elements to speed-up the computation of the structure *parent*. The omitted and facultative procedure *CANONICALIZE_TREE* proposed in [5] transforms the resulting tree into a more compact tree structure, for which each node is a set of pixels of a connected component at a particular level. The tricky part of this algorithm is the processing order \mathcal{R} of the pixels, namely the procedure *SORT*. The sorting defines the type of inclusion in the resulting tree. For example, by sorting the pixels by their grayscale values in an ascending order, darker pixels are visited first, and are then connected by a *parenthood* relationship to an ancestor pixel whose grayscale value is identical or brighter. This results in a tree whose leaves correspond to local minima, and whose global root to the brightest component. This is the min-tree. Similarly, considering a descending order yields

Algorithm 1: “Union-Find”-based computation of a morphological tree.

```

UNION_FIND( $\mathcal{R}$ ) :  $T$ ;
begin
  for all  $p$  do  $zpar(p) \leftarrow \text{undef}$ ;
  for  $i \leftarrow 0$  to  $N - 1$  do
     $p \leftarrow \mathcal{R}[i]$ ;
     $parent(p) \leftarrow p$ ;  $zpar(p) \leftarrow p$ ;
    for all  $n \in \mathcal{N}(p)$  do
      if  $zpar(n) \neq \text{undef}$  then
         $r \leftarrow \text{FIND\_ROOT}(zpar, n)$ ;
        if  $r \neq p$  then
           $parent(r) \leftarrow p$ ;
           $zpar(r) \leftarrow p$ ;
    return  $parent$ ;
FIND_ROOT( $zpar, x$ ) :  $P$ ;
begin
   $y \leftarrow zpar(x)$ ;
  if  $y = x$  then
    return  $x$ ;
  else
     $zpar(x) \leftarrow \text{FIND\_ROOT}(zpar, y)$ ;
    return  $zpar(x)$ ;
COMPUTE_TREE( $u$ ) : Pair(Array[ $P$ ],  $T$ );
begin
   $\mathcal{R} \leftarrow \text{SORT}(u)$ ;
   $parent \leftarrow \text{UNION\_FIND}(\mathcal{R})$ ;
   $\text{CANONICALIZE\_TREE}(parent, u, \mathcal{R})$ ;
  return ( $\mathcal{R}, parent$ );

```

to a max-tree. For the tree of shapes, the order depends on the structural inclusion. We first consider internal components, which are thus placed as leaves in the tree, before processing the external including components. We refer the reader to [5], [6] for more details.

B. Computing the Attributes

A set of attributes is associated to each node, from which it can be determined whether the underlying subtree corresponds to a speckle spot or not. These attributes were derived from the hypotheses *H1-3*. The first hypothesis requires the connected component (i.e. subtree) to be a local maximum or minimum. This property is satisfied by construction in the min-tree and max-tree, but needs to be checked in the case of the tree of shapes. This is done by storing the ordering inclusion type of a node, i.e. increasing or decreasing, in relation to its parent. For *H2* and *H3*, their quantifications are imprecise: regarding the shape of the speckle spots, the discretization and the scatterer density induce some variations from ellipses; with respect to the dimensions of the speckle spots, the lateral resolution depends on the depth of the echo [3].

We modeled these imprecisions by using a fuzzy representation. Indeed, fuzzy sets theory is well adapted to formally describe a vague or imprecise knowledge [7]. In a fuzzy framework, an *Ellipse* is a fuzzy set, defined by a membership function $\mu_{el} : \mathbb{R} \rightarrow [0, 1]$. For an element $x \in \mathbb{R}$, $\mu_{el}(x)$ denotes its degree of membership to the fuzzy set *Ellipse*. In


practice, the element x is defined as the ratio r of the connected component area to its bounding box area. If the component is an ellipse, $r = \pi/4$; if it is a rectangle, $r = 1$. The fuzzy membership function μ_{el} is defined such that it equals 1 when $r \geq \pi/4$. This way, rectangular shapes are not penalized to add flexibility to the detection. For the speckle spot dimensions, we defined two trapezoidal membership functions μ_x and μ_y , the depth and the lateral dimensions, respectively. The membership function μ_x is maximal when the component height belongs to $[\text{FWHM}_x - a, \text{FWHM}_x + a]$, where FWHM_x is the speckle spot height computed using the FWHM of the ACVF. The support of μ_x is set to $[0, 2 \times \text{FWHM}_x]$. The membership function μ_y is similarly defined with the speckle spot width FWHM_y . The parameter a is used in order not to penalize the variations of the speckle spot dimensions along the depth, and can be automatically determined by computing the ACVF at different depths, or by setting it empirically, as we did in our experiments (e.g. $a = \text{FWHM}/4$).

We then merged the membership functions μ_{el} , μ_x , and μ_y , using a fuzzy fusion strategy. This fusion yields to the definition of the fuzzy set *Speckle spot*. Let $\psi: \mathbb{N} \times \dots \times \mathbb{N} \rightarrow [0, 1]$ be a fuzzy fusion operator such that $\psi(\mathcal{S}) = \top(\mu_{el}, \mu_x, \mu_y)$, with \mathcal{S} a set of pixels that belongs to a particular subtree, and \top a t-norm (fuzzy conjunction) [7]. A conjunction is well adapted to our case since a speckle spot must satisfy all the constraints induced by $H2$ and $H3$. In our experiments, we set $\top = \min$. Finally, if the ψ value obtained at a node is above a given threshold λ , then its whole subtree is classified as speckle spot.

IV. SPECKLE REDUCTION


Speckle reduction was performed by pruning the tree of shapes. The image was reconstructed by assigning to the pruned subtrees the grayscale value of their closest non-pruned ancestor, i.e., their including component. We compared our filter to the despeckling methods listed in a recent literature review [8]: a median filter; an ideal low-pass filter; a second-order Butterworth (BW) filter; a wavelet filter, which hard-thresholds LH, HL and HH subbands using a wavelet transform; and the homomorphic versions (abbreviated by H.) of the ideal, Butterworth, and wavelet filters, which carry out the filtering on the logarithmic original images, arguing that the speckle is a multiplicative noise.

The first set of experiments was conducted using a V-1-128 Verasonics research scanner with a 2.5 MHz linear-array transducer. We used a 403GS LE Gammex phantom to scan inclusions (Fig. 2a). Both images were manually segmented into several regions (2 inclusions, and the background), and for each region, its gray level has been set to its mean value. Following [1], [8], we quantitatively evaluated the algorithms using the mean square error (MSE), the signal-to-noise ratio (SNR, computed on small windows all over the image), and the contrast-to-noise ratio (CNR, computed as an absolute normalized difference of the mean intensity between the external and the internal borders). Also, it is imperative not to oversmooth the output image, since it may deteriorate the visibility of thin structures. Thus, we also indicated the gradient norm (Grad norm). Results are depicted on Fig. 2b. Each measure was normalized by the results obtained on the original image. Thus, a MSE inferior than 1, and a SNR and



(a) Inclusions


(b) Results on the inclusions



(c) Original

(d) Median

(e) Ideal



(f) Wavelet

(g) Butterworth

(h) Our approach

Fig. 2. (a),(b): Quantitative results of despeckling methods on two inclusion phantoms. (c)-(h): Qualitative results of despeckling methods on a pig heart.

a CNR greater than 1, indicate some improvements compared to the original image. For the gradient norm, it is preferable to obtain a value lower than 1, but not too small neither. Each method has one parameter (window size, cutoff frequency, and detection threshold for ours), which was optimized to obtain the best results. Our approach showed consistent improvements compared to the original image, which have also been verified on phantoms of different inclusions. Besides, as explained in [1], it is very difficult to evaluate despeckling algorithms since there exists no consensus on a criterion. Therefore, visual assessment is imperative. Results obtained by the best methods on a pig heart are displayed in Figs. 2(e)-(h). Our approach seems to well preserve the contours and small non speckle components, such as specular reflections. This is due to the morphological nature of the method, which operates only on connected components, without any assumption on the image contrast. The output was also successfully flattened in the presence of high density speckles, while not producing oversmoothing, clearly visible for the competitive methods.

V. SPECKLE TRACKING

Speckle tracking was performed in three steps: (1) speckle spots were detected using a morphological tree; (2) their centroids were extracted; (3) a point set registration algorithm matched sets extracted from consecutive images. This pipeline differs from other speckle tracking algorithms, in that speckle

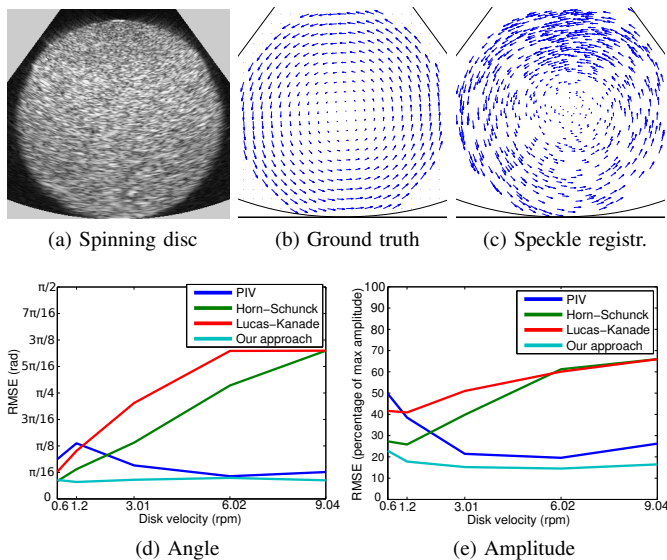


Fig. 3. Top: Vector flows obtained on a spinning disc at $\omega = 6.02$ rpm. Bottom: RMSE in angle and amplitude for the different angular velocities.

spots were explicitly extracted. Thereby the tracking algorithm was employed only on the most relevant data. We used the registration algorithm proposed in [9], which represents the point sets by Gaussian mixture models, and then minimized the alignment of the two distributions according to the L2 distance. The non-rigid transformations were operated by a thin-plate spline model. In order to limit the number of points to be registered, only the most contrasted bright speckle spots (700) were detected. Detecting bright spots could be done using a tree of shapes by excluding the local minima. However, a max-tree provides the same detection results, with the benefit of computing the morphological tree more efficiently. Also, considering only the most contrasted 700 spots reduced the registration algorithm complexity while making it more robust. The contrast was computed as an attribute in the max-tree, and was defined as the absolute difference between a component and its parent component. We compared our approach with the Horn-Schunck (H-S) and Lucas-Kanade (L-K) optical flow methods [10], and with the particle image velocimetry (PIV) method [11]. The H-S method formulates the optical flow estimation as a global minimization procedure integrating a smoothing constraint. The L-K method assumes a constant flow within small local neighbourhood windows, which leads to an overdetermined problem that is solved in the least squares criterion. The PIV method estimates the velocity field by computing multi-scale regularized cross-correlation windows.

Experiments were conducted using a Verasonics research scanner with a 2.5 MHz phased-array transducer. 25 images were acquired per second. The transducer imaged an ultrasound disc phantom spinning at angular velocities ranging from 0.6 rpm to 9.04 rpm. Figure 3 (top) depicts the phantom, the ground truth velocity map at $\omega = 6.02$ rpm, and the tracking results obtained by our approach. The registration obtained on the detected speckle spots is illustrated in Fig. 3c, and can thus be interpolated on a dense grid. Quantitative results obtained by the optical flow methods, the PIV, and our approach, are presented in Figs. 3d and 3e. Evaluation was

quantified using the angle (in radians) and amplitude (in percentage of the maximum amplitude) root MSE (RMSE) of the flow velocity. For each angular velocity, results were averaged over 15 consecutive images. Our approach performed well for every angular velocity. Similar results were also observed using ultrafast imaging (circular waves, 11-angle compounding, 100 fps after compounding). Because of the intensity conservation hypothesis and the first order approximations, errors observed for the optical flow methods increased with angular velocity. PIV showed bad performances when the angular velocity was too small, due to the shift towards integer values inherent in this method, which biases the estimation especially when the motion is slow. Our approach did not suffer from any of these limitations, since the extracted centroids were real-valued, and there was no assumption made on the formation of the image.

VI. CONCLUSION

We proposed a morphological tree to detect speckle spots in ultrasound images, in which the spots are detected using structural characteristics. This inclusion representation is well adapted to reduce speckle as it showed nice properties compared to state-of-the-art methods: it is contrast independent, uses only structural information, and preserves thin components and contours. We also applied this detection method for speckle tracking. Paired up with a point set registration algorithm, the extracted centroids were accurate enough to handle small as well as large displacements. As a perspective, a morphological tree might be used to detect more than just speckle spots, e.g. specular reflections, shadowing artifacts, and abnormal tissues.

REFERENCES

- [1] C. P. Loizou, C. S. Pattichis, C. I. Christodoulou, R. S. Istepanian, M. Pantziaris, and A. Nicolaides, "Comparative evaluation of despeckle filtering in ultrasound imaging of the carotid artery," *IEEE TUFFC*, vol. 52, no. 10, pp. 1653–1669, 2005.
- [2] P. Monasse and F. Guichard, "Fast computation of a contrast-invariant image representation," *IEEE TIP*, vol. 9, no. 5, pp. 860–872, 2000.
- [3] R. Wagner, S. Smith, J. Sandrik, and H. Lopez, "Statistics of speckle in ultrasound B-scans," *IEEE Transactions on Sonics and Ultrasonics*, vol. 30, no. 3, pp. 156–163, 1983.
- [4] G. R. Bashford and O. T. von Ramm, "Speckle structure in three dimensions," *The Journal of the Acoustical Society of America*, vol. 98, no. 1, pp. 35–42, 1995.
- [5] C. Berger, T. Géraud, R. Levillain, N. Widynski, A. Baillard, and E. Bertin, "Effective component tree computation with application to pattern recognition in astronomical imaging," in *IEEE ICIP*, vol. 4, 2007, pp. IV–41.
- [6] T. Géraud, E. Carlinet, S. Crozet, and L. Najman, "A quasi-linear algorithm to compute the tree of shapes of nD images," in *ISMM*, 2013, pp. 98–110.
- [7] D. Dubois and H. M. Prade, *Fuzzy sets and systems: Theory and applications*. Academic Press, 1980.
- [8] J. L. Mateo and A. Fernández-Caballero, "Finding out general tendencies in speckle noise reduction in ultrasound images," *Expert Systems with Applications*, vol. 36, no. 4, pp. 7786–7797, 2009.
- [9] B. Jian and B. C. Vemuri, "Robust point set registration using gaussian mixture models," *IEEE TPAMI*, vol. 33, no. 8, pp. 1633–1645, 2011.
- [10] D. Fleet and Y. Weiss, "Optical flow estimation," in *Handbook of Mathematical Models in Computer Vision*. Springer, 2006, pp. 237–257.
- [11] D. Garcia, "A fast all-in-one method for automated post-processing of PIV data," *Experiments in fluids*, vol. 50, no. 5, pp. 1247–1259, 2011.