Pattern Recognition 47 (2014) 2467–2480

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Detection of masses and architectural distortions in digital breast tomosynthesis images using fuzzy and *a contrario* approaches

Giovanni Palma [a,*], Isabelle Bloch [b], Serge Muller [a]

[a] GE Healthcare, 283, rue de la minière, 78530 Buc, France
[b] Institut-Telecom, Telecom ParisTech, CNRS LTCI, 46, rue Barrault, 75013 Paris, France

ARTICLE INFO

Article history:
Received 11 October 2012
Received in revised form
9 December 2013
Accepted 22 January 2014
Available online 1 February 2014

Keywords:
Digital breast tomosynthesis
Computer aided detection
Masses
Architectural distortions
Fuzzy connected filter
A contrario detection
Classification
Support vector machine

ABSTRACT

Digital breast tomosynthesis (DBT) is a new 3D imaging technique, which overcomes some limitations of traditional digital mammography. Its development induces an increased amount of data to be processed, thus calling for a computer aided detection system to help the radiologist. Towards this aim, this paper focuses on the detection of masses and architectural distortions in DBT images. A complete detection scheme is proposed, consisting of two parts, called channels, each dedicated to one type of lesions, which are then merged in a final decision step, thus handling correctly the potential overlap between the two types of lesions. The first detection channel exploits the dense kernel nature of masses and the intrinsic imprecision of their attributes in a fuzzy approach. The second detection channel models the convergence characteristics of architectural distortions in an *a contrario* approach. The experimental results on 101 breasts, including 53 lesions, demonstrate the usefulness of the proposed approach, which leads to a high sensitivity with a reduced number of false positives, and compares favorably to existing approaches.

1. Introduction

Digital breast tomosynthesis (DBT) is a new three-dimensional (3D) imaging technique aiming at overcoming some limitations of mammography [1]. It has the potential of improving the visibility of breast structures by reducing overlap of tissues. Therefore, the detectability of lesions is potentially increased while false positives (FPs) due to tissues superimposition can be more easily discarded. This comes at the expense of an increased amount of data to be reviewed by the radiologist. Therefore the need for an automatic detection/characterization system, also known as computer aided detection (CAD), dedicated to lesions contained in this new kind of data is increasing, since it may help the radiologist to achieve detection tasks in a reasonable amount of time, while keeping or increasing his sensitivity. Towards this aim, we propose in this paper a complete detection scheme for automated detection of masses and architectural distortions, which are suggestive of malignancy. Fig. 1(a) illustrates the former type of lesions. It is characterized by a dense kernel and ill defined boundaries. In this particular example there are also some spicules. Fig. 1(b) represents the latter, which is characterized by the lack of dense kernel and a strong convergence pattern.

The aim of this paper is not to identify the clinical type of lesions, but detect lesions whatever their type. All lesions we consider have one of the two different appearances in the images. Therefore we design a detection method for each of these broad classes of appearances, regardless of the exact clinical type of the lesion. Since there is no crisp transition between the two appearance classes, the results are finally grouped to provide better overall detection results. While a third type of finding exists (calcification clusters), it is not considered in this paper. Actually, we rely only on commonalities between masses and architectural distortions in order to detect border line findings like strongly spiculated masses. Although this reasoning makes sense for these two types of lesions, it does not for calcification clusters since they do not share characteristics with them. Therefore no gain in overall detection performance can be expected compared to adding an external approach like for instance [2] to our method. For the same reason, CAD systems in the literature are usually dedicated to either masses or calcifications.

Although DBT aims at offering a better characterization of radiological findings, the variability of these structures is still large. Thus, as it is the case for standard mammography, the detection task of pathological patterns remains challenging. Furthermore, DBT being a quite recent imaging technique, the literature does not provide a comprehensive solution to the CAD

* Corresponding author. Tel.: +33 1 30 70 96 17; fax: +33 1 30 70 94 35.
E-mail addresses: giovanni.palma@ge.com (G. Palma), isabelle.bloch@telecom-paristech.fr (I. Bloch), serge.muller@ge.com (S. Muller).

0031-3203/$ - see front matter © 2014 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.patcog.2014.01.009

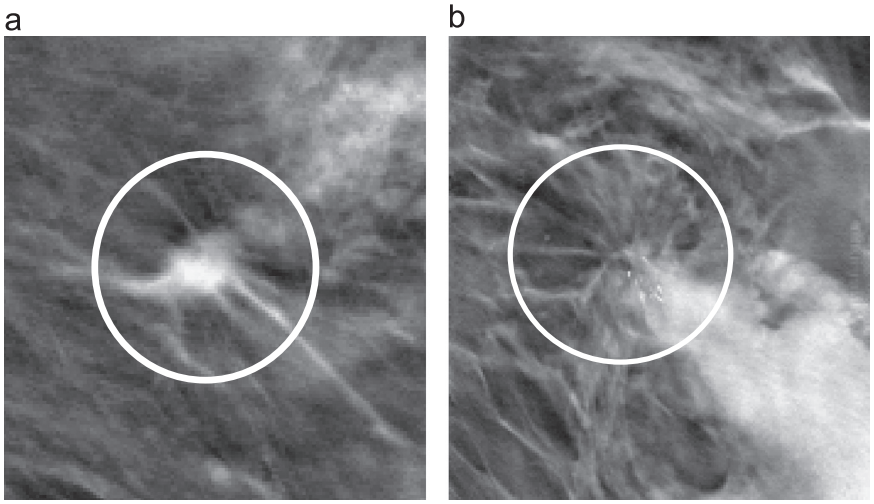

Fig. 1. An example of mass (a) and architectural distortion (b) in DBT.

problem, especially for the case of architectural distortions, which has not been fully addressed so far.

In DBT-CAD, sensitivity can be reported as the ratio of findings that are detected. This can be done either on a DBT volume or on a breast basis. In the first case, each lesion in each DBT volume is considered as a finding to be detected. In the second case, a given lesion contained in a breast that has been acquired from several views (i.e. several DBT volumes) only needs to be detected a single volume. In the literature, the first one is commonly used, mainly because early DBT exams usually contain only one view. However, sometimes both measurements are used [3]. For the sake of clarity, we will always use the first method in this paper when presenting existing results or the performance of our approach. Additionally, specificity will always be given as the number of false positives per DBT-volume.

Preliminary investigations were based on a detection of masses performed within the projections to be used for the reconstruction of the volume [4]. The same authors also proposed to work directly within the DBT volume [5]. While these developments demonstrated the feasibility of a CAD system, specificity remained high (2.2 FP per breast volume at 85% sensitivity).

Other groups worked on both projections and volume based approaches. Reiser et al. [6] proposed to detect masses independently within each projection and to recombine the detection output in the 3D space using visibility angular range of the findings. They also proposed a processing of the volume by a radial gradient filter combined with a maximum intensity projection and a top hat [7]. This processing produces a large number of false positives, therefore it is only suitable as a preliminary detection step of a whole CAD chain.

A fuzzy logic based processing chain has also been proposed [8]. It relies on an independent detection of the suspicious areas within the projections, followed by a segmentation stage using a fuzzy contour framework [9], which allows handling imprecision and uncertainty regarding these steps until the decision making stage using fuzzy decision trees [10]. However, this work mainly focuses on the information aggregation part and lacks a robust initial detection of findings.

More recently a projection and a volume based CAD system were compared by Chan et al. [11]. A hybrid method that combines both approaches was also introduced. The authors used quite a large database (69 malignant lesions) for the evaluation and concluded that their 2D only approach performed worse than the hybrid one. While this work focuses only on mass detection, it will be considered as the reference method in this paper and a comparison with our approach will be provided later on, mainly because of the database size and because it provides the best performance in the literature (1.61 FP per breast volume at 90% sensitivity).

Another kind of approach uses information theory [12,13] to reduce the number of false positives. While results are promising, the false positive rate still remains high (2.4 FP per breast volume at 90% sensitivity) compared to the method of Chan et al. [11]. This kind of method will be discussed in this paper as well.

Finally, a more recent work has been proposed [3]. The idea is to use a CAD system trained on 2D mammograms on tomosynthesis slabs (aggregation of slices). Using the same evaluation methodology as the other papers, the proposed method gives performance slightly better than the approach of Singh et al. [12]. However, the use of slabs and the poor slice inter-spacing have a negative impact on the localization of the lesion in the volume.

From a technical perspective, designing an efficient CAD system is quite a challenging task mainly because of the difficulty to mathematically define the findings of interest in a way that captures their variability. This results in the difficulty to derive robust operators for their detection/characterization. Actually, existing systems usually rely on two steps. First the detection of potential findings and then false positive reduction. The former should be as sensible as possible in order not to miss cancers. The latter has to drastically reduce the number of false detections. However, if the early detection has a poor specificity, the second stage will likely not reduce it to an acceptable level. In this paper, we introduce a method to cope with all these constraints altogether.

In our work, we propose a detection scheme that addresses both masses and architectural distortion with similar performance to state of the art mass only DBT CAD systems. Since these findings have different appearances and characteristics in the images, we propose an original scheme, composed of two channels, each one being dedicated to one type of lesions. The results provided by these two channels are then aggregated to reach the final decision on the detection. In the first channel, a fuzzy approach is implemented to detect masses. The idea is to model the imprecision on the contours using fuzzy logic, which is suitable to propagate/handle this type of imperfection through the detection process. In the second channel, convergence regions are detected using an *a contrario* approach. Here the idea is to statistically define the content of a healthy breast in order to detect abnormalities.

An overview of the proposed approach is given in Section 2. Then the detection procedures in the two channels are described

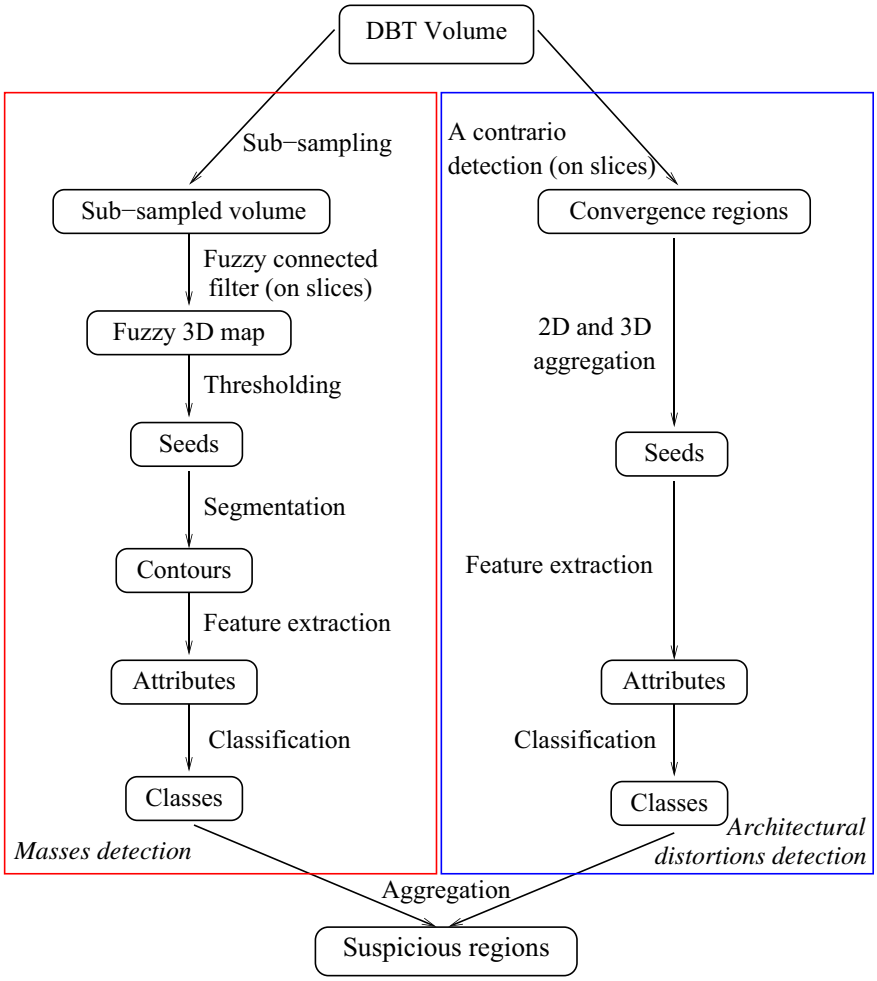

Fig. 2. Global scheme of the two-channels detection scheme for masses and architectural distortions.

in Sections 3 and 4. In Section 5 we discuss the aggregation step and the obtained results.

2. Method overview

In this section we provide a general overview of the proposed approach, and we describe the database used for the experiments and the evaluation.

2.1. Proposed approach: a two-channel scheme

One of the original features of the proposed approach is to deal with two types of lesions separately. This allows taking into account the specificities of each type of lesion before aggregating the results. The overall scheme, which is applied directly on the reconstructed tomosynthesis volume, is illustrated in Fig. 2. The detection of masses (e.g. Fig. 1(a)) relies on the recognition of dense kernels, while the detection of architectural distortions (e.g. Fig. 1(b)) is based on the identification of convergence regions. Indeed, these two radiological signs are likely to represent the two types of lesions to be detected. It should be noted that another channel could be added for detecting microcalcifications, for instance using the method described by Bernard et al. [2], thus leading to a complete system for suspicious findings detection.

The first channel, which is dedicated to dense kernels, relies on the modeling within a fuzzy set framework of both lesion imprecision and variability in the image. The tomographic planes of the acquired volume are first processed one after the other. Focal densities are detected using an original fuzzy connected filter derived from the one introduced by Palma et al. [14] and involving size, compactness and contrast attributes. A fuzzy map of potentially suspicious regions is then computed, using 3D pseudo-connected component labeling based on a maximum shape variability criterion. This allows disconnecting some components and ensuring that each region corresponds to only one potential lesion and does not aggregate two distinct structures together. Then each suspicious region is segmented. The decision on whether the suspicious region is actually a lesion or not is made from a support vector machine (SVM) classification [15] based on specific attributes (such as compactness, size, statistical analysis of the neighborhood). This final segmentation and classification step aims at reducing the number of false alarms.

The second channel aims at detecting suspicious convergence regions, indicating potential architectural distortions. The characteristics of these regions led us to propose an *a contrario* modeling [16] of the problem. This approach consists in statistically defining the content of healthy breast volumes allowing detecting patterns in real images that are unlikely to appear in such data. This is performed on each tomographic plane, before aggregating the results in the 3D volume. The final step is again a classification based on specific features (e.g. orientations analysis in a neighborhood of the finding).

Once each channel has been applied independently, a final fusion step is performed. Since the channels have dedicated and distinct objectives, a direct disjunctive fusion can be applied and a

region in the volume is considered as suspicious if it has been detected by at least one channel. It should be noted that the two classes of lesions are not perfectly separated. For instance a highly spiculated mass may be detected by both convergence detection and dense kernel channels. Therefore, although not sophisticated, the fusion actually improves the performance of the overall chain, which is one contribution of this paper.

While these two channels rely on existing techniques that have recently been introduced, their refinement and implementation in a full detection scheme demonstrate their validity for a clinical task.

2.2. Image database

The database used in our experiments is composed of 101 DBT acquisitions (1 breast, 1 MLO view for each acquisition), 53 with a biopsy proven cancer and 48 with no pathology. Localization of lesions in the DBT volumes has been performed according to radiologist's report defining the exact biopsy target (and therefore the cancer if malignancy is proven). Contours of the lesions have been drawn according to radiologist information such as lesion size and validated by two clinical experts (cases were reviewed and refined by the reviewers together). The repartition of the 53 lesions is as follows:

- 7 irregular masses,
- 4 lobulated masses,
- 39 spiculated masses,
- 3 architectural distortions.

DBT volumes were reconstructed using iterative techniques [17] from low dose projections acquired over an angular range of 40°. Slice inter-spacing and in-plane spatial sampling were set to 1 mm and 0.1×0.1 mm, respectively.

In order to derive clinically relevant figures, the images used for the computation of the specificity index do not contain any suspicious region detected by radiologists, while the images used for the computation of the sensitivity index contain at least one lesion detected by radiologists, with biopsy proven malignancy.

The density detection channel was evaluated using all irregular and lobulated lesions, as well as the less visually spiculated masses (i.e. a total of 40 lesions). The remaining lesions, including the highly spiculated lesions and the architectural distortions, were used to evaluate the second channel. The choice to split the spiculated lesions into two pools was motivated by the fact that the most characteristic feature of such lesions is sometimes their stellate pattern rather than their density, and this feature is mainly exploited by the second channel. From an experimental standpoint, this choice also leads to a more accurate evaluation of the convergence detection channel, with 13 lesions, even if this number is still too low to lead to definite conclusions about a clinical evaluation. Still, we believe that this provides a proof of concept and an evaluation of the interest of the proposed approach.

3. Detection of masses

The first channel of our detection scheme relies on the detection and classification of dense kernels. In this section, we focus on the description of the detection techniques implemented in the first stage, the segmentation methods and the potential lesions classification.

3.1. Dense kernels detection

The detection of dense kernels contained in the image is performed on an under-sampled version of the original volume. This step allows speeding up the computation time with limited impact on the detection performance because of the scale of target structures. Actually, masses seen in DBT are usually quite large (from several millimeters to several centimeters) in comparison to other findings like microcalcifications, which require a full resolution to be seen and detected.

3.1.1. Fuzzy connected filters

Potential masses are usually detected either using a gradient analysis [11] or with a difference of Gaussian's [12]. These two methods work on circular regions, which may not reflect the shape of all lesions. Connected filters may help to relax this assumption because the shape of potential lesions is given by the signal itself (connected component resulting from image thresholding). Furthermore, this kind of filters has recently been extended to the fuzzy set framework [18,14] making them suitable to handle uncertainty of lesions on both the shape of connected components and the output magnitude with regard to selected criteria. The assumption that motivates the use of such filters as a detection step is that dense kernels can roughly be retrieved by multi-thresholding tomosynthesis images. Actually, if the object we try to detect is denser than its surrounding, it can be retrieved by thresholding and extracting a connected component. The proposed assumption holds because of the 3D nature of the image: DBT suppresses a large part of tissue superimposition compared to standard mammography.

More specifically, fuzzy connected filtering relies on two features. First the concept of image is replaced by the one of fuzzy image. For a given image I, for each pixel p in the image domain Ω, the associated value $I(p)$ is no longer a crisp value belonging to the set of gray-levels $\mathcal{G}$ but a fuzzy quantity (columns in Fig. 3). This fuzzy set is defined on the domain of possible gray-levels where a membership degree between 0 and 1 is associated with each possible value belonging to $\mathcal{G}$. Fuzzy umbra images are defined as

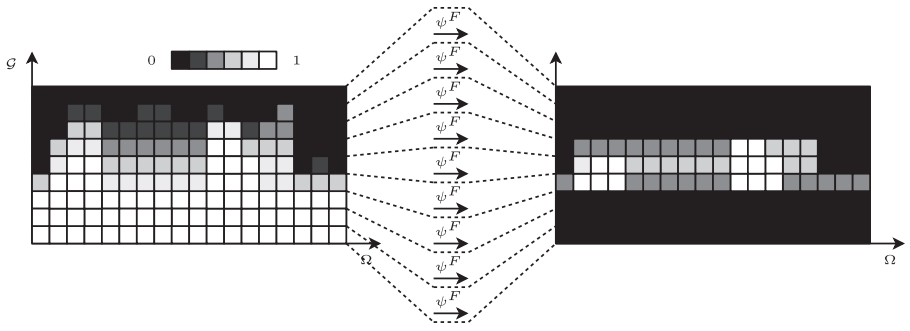

Fig. 3. Example of a fuzzy connected filter δ. On the left: a fuzzy umbra image (F). On the right: the filtered fuzzy image. For every gray level $g \in \mathcal{G}$, a fuzzy set $F(*,g)$ is extracted and filtered using ψ^F.

fuzzy images where the fuzzy quantities associated with every pixels are decreasing with respect to the gray value (see left image in Fig. 3). As it will be shown later, this representation will enable us to consider fuzzy thresholding of the image. Let us call $\mathcal{F}$ the set containing all the possible fuzzy images. This set can be seen as the set of all the fuzzy sets defined on $\Omega \times \mathcal{G}$. Using this notation, for a given pixel $p \in \Omega$, for a given gray-level $g \in \mathcal{G}$, a fuzzy umbra image $F \in \mathcal{F}$ is interpreted as follows: $F(p,g)$ is the degree to which the image has a gray-level greater than or equal to g at pixel p. The second aspect of this framework concerns the way images are processed. A fuzzy filter is defined using a set of operators ψ^F to be applied on the sets $f = F(*,g)$ extracted for each gray level g. These fuzzy sets f are fuzzy sets defined on the image domain Ω extracted from the fuzzy image F for a given gray level g ($*$ denoting all the elements of Ω). It corresponds to rows of fuzzy images in Fig. 3. The membership degree of f for each pixel p is the degree of the fuzzy quantity associated with the considered gray value ($F(p,g)$). In the case of fuzzy umbra images, for a given gray-level, the corresponding fuzzy set can be seen as a fuzzy thresholding of the input image. Fig. 3 illustrates the filtering of a fuzzy image. The set of all the fuzzy sets on Ω is denoted by $\mathcal{S}$. In our CAD application, the filter we use on the various gray levels extracts all the fuzzy connected components $\mathcal{H}(f)$ [19–22] from the previous sets f and keeps only the ones that verify simultaneously size, compacity and contrast-based criteria. This filter is defined in the below equation

$$\forall F \in \mathcal{F},\ \forall f \in \mathcal{S},\ \forall p \in \Omega,$$

$$\begin{aligned}\psi^F(f)(p) = &\max_{h \in \mathcal{H}(f)} (\min(h(p), t_{u_1}(fcard(h)))) \\ &\top \max_{h \in \mathcal{H}(f)} (\min(h(p), r_{u_2}(fcomp(h)))) \\ &\top \max_{h \in \mathcal{H}(f)} (\min(h(p), r_{u_3}(frose(F,h))))\end{aligned} \quad (1)$$

where $\top$ is a t-norm, which is a fuzzy conjunction (fuzzy extension of a logical and), t_u and r_u are the ramp functions parameterized by u, *fcard* and *fcomp* are the fuzzy cardinality (see Eq. (2)) and fuzzy compacity (see Eq. (3)), respectively [23,18] and *frose* the operator of Eq. (4). Let us note that this last criterion, which represents a contrast that takes into account the size of the object, is one contribution of this work

$$fcard(h) = \sum_{p \in \Omega} h(p) \quad (2)$$

$$fcomp(h) = \frac{fcard(h)}{(\sum_{p \in \Omega} |\nabla h(p)|)^2} \quad (3)$$

$$\forall F \in \mathcal{F},\ \forall f \in \mathcal{S},\ \forall \nu \in \mathcal{S},$$
$$frose(F,f) = \frac{fctrast(F,f)}{fmean(F, D_\nu(f) \cap \bar{f})} \sqrt{fcard(f) fmean(F, D_\nu(f) \cap \bar{f})} \quad (4)$$

In Eq. (4), $D_\nu(f)$ corresponds to the fuzzy dilation of a set f using a fuzzy structuring element ν (a disk in our case) as defined by Bloch and Maitre [24], *fmean* denotes the fuzzy mean (i.e. weighted mean of F by $D_\nu(f) \cap \bar{f}$), and *fctrast* denotes the contrast, computed as the difference of the means inside and outside the component (where the outside is restricted to $D_\nu(f) \cap \bar{f}$).

The filter ψ^F can be interpreted as follows: for each point p of the image domain, the resulting degree corresponds to the conjunction (expressed by the t-norm $\top$) of the degrees associated with the criteria results computed on the various fuzzy connected components $h \in \mathcal{H}(f)$ that have been extracted for a given fuzzy set f. This means that the degree at point p will be high if it belongs to a fuzzy connected component that verifies the criteria of size, compacity and detectability.

As said previously, the filter we have just introduced is dedicated to the processing of fuzzy sets contained in a fuzzy image and extracted for the various gray levels it encodes. Defining a filter δ that processes all the fuzzy sets $F(*,g)$ extracted from a fuzzy image F (see Fig. 3) is then straightforward as shown in the below equation

$$\forall F \in \mathcal{F},\ \forall p \in \Omega,\ \forall g \in \mathcal{G},$$
$$\delta(F)(p,g) = \psi^F(F(*,g))(p) \quad (5)$$

The detection filter that is applied to the volume slices I is then expressed as $\phi(I) = agg(\delta(um(I)))$, with $\forall p \in \Omega,\ \forall g \in \mathcal{G},\ \delta(um(I))(p,g) = \psi^{um(I)}(um(I)(*,g))(p)$. The *um* operator enables to consider the original image as an umbra image. It is responsible for the modeling of the imperfections like noise arising from the image [25]. The *agg* operator as originally introduced by Palma et al. [18] allows aggregating the filtering results associated with each gray level (e.g. $agg = \max$ for a disjunctive aggregation).

Let us note that the operator ψ^F is similar to the one originally introduced by Palma et al. [18]. The only difference resides in the measure that is actually used to assess the contrast of a connected component. In the present work, we introduce a new measure, which is a fuzzy version of the signal to noise ratio expressed in the *frose* modeling where the size of the object is taken into account [26]. This criterion was originally defined to assess the detectability of a circular structure of known size in an image corrupted with Poisson noise. By extension, we can use it to detect findings. While the content of DBT planes does not verify the noise assumption, it turned out experimentally that the measure was more suitable than the regular fuzzy contrast.

3.1.2. Various levels of fuzziness

The previously proposed framework allows taking into account imperfection of the data at various stages. From an implementation standpoint, we could use a simplified version of the fuzzy umbra image framework: the *um* operator would actually provide regular umbra images. The proposed filter would then take only advantage of the fuzzy detection output and discards the uncertainty coming from the gray levels. Considerations on computation time could motivate this choice. Actually, while efficient generic algorithms, based on tree update strategies, have already been proposed in the literature [14], processing a crisp umbra image allows us to make simplifications resulting in computation speed up: a single tree representation of the nested connected component is enough to filter the whole image. From a detection standpoint this would obviously lead to a coarser output of the filter. Nonetheless this could be acceptable since in the detection stage of our processing chain we want a high sensitivity even if the specificity is not that good. In this context, the criteria we used are rather flexible and a fine study of the imprecision is not mandatory: the characterization stage is actually performed later.

Practically, we can note that crisp umbra images lead to a rather robust output [18] even if the detection performance could benefit from a comprehensive use of the fuzzy image framework [25] or by modeling some DBT volume characteristics (e.g. reconstruction artifacts).

3.1.3. Illustration

Fig. 4 presents the kind of results we get using the previous filter on a DBT volume. The obtained fuzzy detection map clearly highlights the lesion, which appears as the brightest component. It contains only a limited number of other components, with a much lower degree of detection, which will allow us to recognize them as false positives in the subsequent steps.

3.2. Marker extraction from the fuzzy detection map

Once all the tomographic planes have been filtered, a marker extraction stage is performed in order to better characterize the

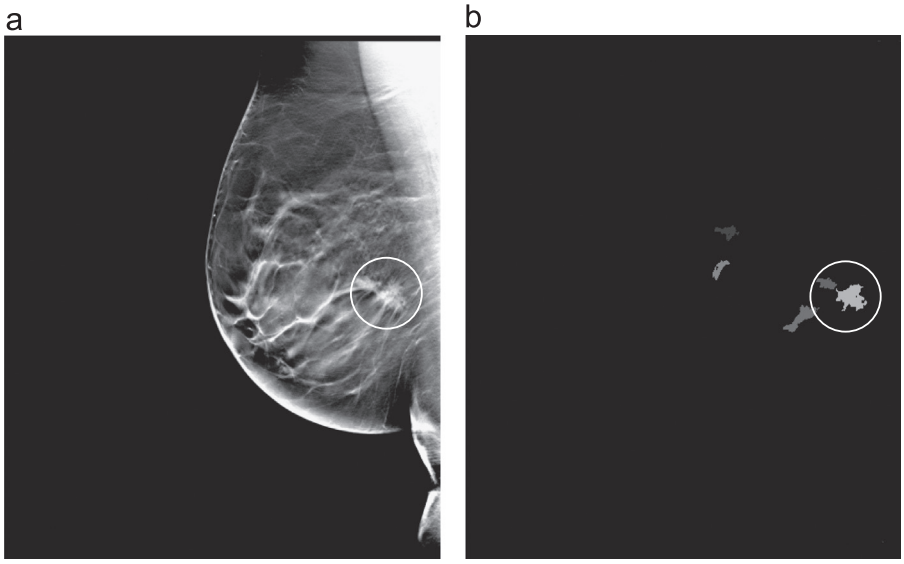

Fig. 4. Illustration of the fuzzy connected filtering of a DBT slice containing a spiculated lesion. (a) Plane containing the lesion (circled in white). (b) Fuzzy detection map (high gray levels correspond to high membership values).

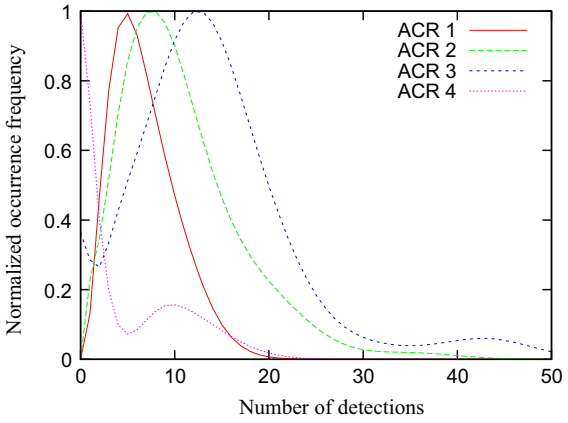

Fig. 5. Normalized histograms for the number of detections with respect to the breast density (ACR classification is used: from 1 for least dense to 4 for the densest breasts).

detected regions and thus reduce the amount of false positives. Markers are sets of connected voxels in the volume domain.

3.2.1. Adaptive thresholding

An adaptive thresholding was implemented to extract the markers from the fuzzy detection map. Actually the texture can change from one breast to another, thus a single threshold does not allow processing in a suitable way (i.e. with a reasonable specificity) a large set of exams: some volumes produce only one or two markers while for some other ones we can reach twenty or more markers.

In order to illustrate the dependency between the number of detections and the composition of the breast, we can refer to Fig. 5, which presents the distribution, estimated on 122 normal cases not used for assessing the algorithms, of the amount of unlimited detections (i.e. with a sufficiently large threshold), for various breast densities. The classification of the density is done using the scale ranging from 1 to 4 proposed by the *American College of Radiology* [27]. The density was assessed by experienced radiologists on the DBT volumes. Let us note that these histograms are highly smoothed. Furthermore, since densities 1 and 4 are less likely to appear than others, their histograms are computed using only a small number of cases, and thus their accuracy is limited. In this illustration, we can see that denser breasts seem to produce a larger amount of false positives. Thus it seems reasonable to adjust the threshold according to the data at the marker detection stage.

Therefore, the proposed strategy consists in automatically restricting the threshold while the amount of detections in the volume is larger than what we are willing to accept. This also enables to tune in a more intuitive and more efficient way the performance of the marker extraction stage. This number of detections corresponds to the number of markers to be presented later on to further processing steps.

Again, let us recall that although the histograms of Fig. 5 help us in defining an adaptive thresholding mechanism, we need to be careful before drawing any other conclusions from its content, mainly because of the small amount of data for some classes.

3.2.2. 3D aggregation

In order to take advantage of the 3D information, connected components resulting from the former adaptive thresholding are gathered in 3D. This gathering is done using the connectivity between the detection results obtained from consecutive planes, in combination to a similarity measure. This similarity is measured by computing the shape variation (see Eq. (6)) between several markers $\{C_z\}$, each C_z being the set of connected points from the same plane z contained in the marker. The cardinality of a marker is denoted by $|\cdot|$.

$$\text{variation}(\{C_z\}) = \frac{|\cup_z C_z|}{\max_z |C_z|} \tag{6}$$

Concretely, the aggregation of the tagged regions in the volume is done as follows: we consider every pair of 2D markers, which have been extracted from the various planes that are connected together on the Z-axis. These pairs are aggregated if the resulting 3D marker has a variation value less than or equal to a given threshold. This ensures not to consider markers with a high dissimilarity in the remaining processing. This threshold has been set once to a restrictive value in order to avoid any meaningless aggregation. Nonetheless this may result in a non-aggregation of markers that should actually be put together. While being non-optimal, such a behavior is acceptable since if a single lesion is split into two separate components, it will not be discarded but

rather processed several times later on. Thus, the choice of this threshold is rather flexible.

3.2.3. Intermediate performance

This marker detection step has been assessed using cross-validation techniques on the database containing the lobulated, irregular and spiculated lesions that were not used for the convergence detection assessment. For this purpose, 40 malignant lesions were used, as explained in Section 2.2. Fig. 6 illustrates the various sensitivities associated with the false positive rates that were obtained by varying the maximum number of allowed markers as explained previously.

For instance, at this stage, a sensitivity of 82% (rate of actual malignant lesions detected) is reached for an average number of 3.83 false positives (number of begin areas detected) per volume. Nonetheless, in order to keep as much potential findings as possible, a sensitivity of 100% will be used for the remaining steps even if the false positives rate reaches 6.8 markers per volume (see Fig. 6). The remaining segmentation, feature extraction and classification steps will then aim at lowering this rate.

3.3. Segmentation

Each marker produced by the former stage is considered as a set of voxels inside the volume. Since objects in DBT volumes usually suffer from a distortion in the Z direction, potential lesions are segmented in 2D in the (X, Y) plane parallel to the detector. The plane to be considered for this segmentation step is the most representative one. In order to determine it, a contrast measure is computed for each plane that intersects the marker by subtracting the mean inside and outside the marker. Assuming that the highest contrast is obtained in the focal slice of the lesion, the most contrasted plane is considered as the most representative one. Let us remark that even though segmentation is performed in 2D, 3D information is used through the 3D aggregation step. Furthermore, since this last step tends to over-segment the 3D markers, the set of slices to be considered is usually small and their content is rather similar, which implies that no meaningful information is lost by selecting the most representative plane.

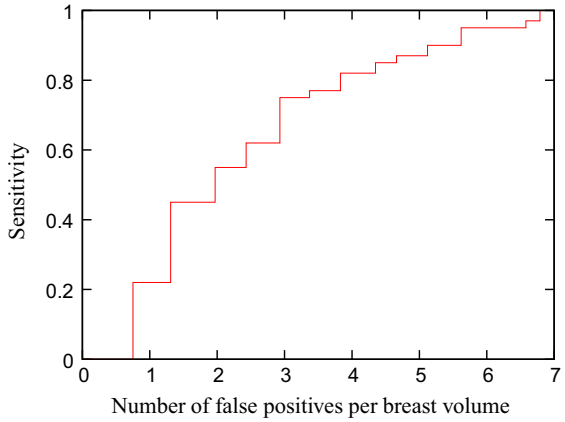

Fig. 6. Performance of the marker stage used for the detection of dense kernels.

We assessed our CAD system using two different segmentation methods. The first one is based on a multi-thresholding strategy while the second one relies on a minimum cost path modeling [28,29].

The multi-thresholding method relies on the same assumption as the one we used for the connected filter implementation (dense kernels can be segmented by image thresholding). For a given seed, the connected components containing this seed and coming from all the possible thresholds of the image are considered. Each component is a potential contour and a set of fuzzy criteria (similar to the ones used for the proposed fuzzy connected filter) are computed and aggregated together. The final contour is the one associated with the highest membership degree.

The second segmentation method relies on the modeling of the segmentation problem as a minimal path search, using an extension of the approach of Timp and Karssemeijer [28]. First the image is converted into the polar domain (the first axis represents the distance to the finding center and the second one represents an angle ranging from 0 to 2π that covers the whole surrounding of the lesion). Once the image is converted in this new representation, some features, like gradient orientations or optimal gray values, are extracted from the contour and used to derive a cost matrix. The segmentation problem is then reformulated as finding a path that crosses all the columns of the cost matrix, corresponding to the different angles, and that has a minimal cumulative cost [29].

Although these segmentation methods are suitable to provide fuzzy contours [8,10,29], we used the crisp versions in the experiments presented in this paper. Actually, for computation time considerations, we chose not to handle imprecision arising from the detections.

Table 1
Features used for the classification of dense kernels.

Criterion	Description
Homogeneity	Variance within the ROI that contains the marker
Entropy	Entropy within the ROI
Contrast	Contrast of the structure (difference between the means inside and outside the contour)
RGI	Radial gradient index on the contour [30]
RGI inside the contour	Likewise but inside the contour
RGI within the region of interest	Likewise but inside the ROI
Mean gradient on the contour	Mean of the gradient, which is computed using Gaussian derivatives on the image
Entropy of orientations inside the contour	Statistical measure computed on the (non-)weighted histogram of the orientation inside the contour
Entropy of orientations on the contour	Likewise but the histogram is computed on the contour boundary
Entropy of orientations within the ROI	Likewise but the histogram is computed using the whole neighborhood of the lesion, i.e. the ROI
Kurtosis, variance, mean, entropy, skewness of relative orientations inside the contour	Statistical measures computed inside the contour on the histogram of the orientations relative to the center of gravity of the contour
Kurtosis, variance, mean, entropy, skewness of relative orientation on the contour	Likewise but the histogram is computed using the elements present on the contour boundary
Kurtosis, variance, mean, entropy, skewness of relative orientation within the ROI	Likewise but the histogram is computed with all the values of the ROI containing the lesion
Minimal convergence probability	Measure derived from the *a contrario* modeling proposed in Section 4 ($\min\{P[Z_{c,r} \geq z_{c,r}], r \leq R_{\max}, c \in \text{contour}\}$, with $z_{c,r}$ the instances of $Z_{c,r}$ for the considered image)

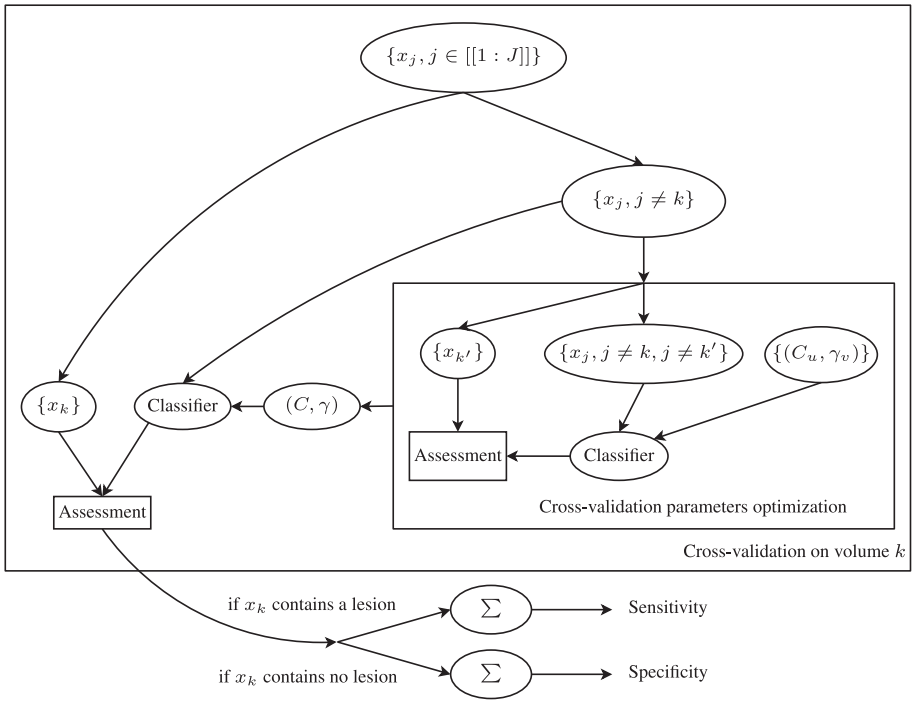

Fig. 7. Cross-validation of the dense kernel validation channel (J is the total number of volumes).

3.4. Features

In order to characterize the contours produced by the former step, several measures are computed for each of them. These features enable to use a classifier as explained later.

Among these features, some of them use only the contour's shape, while others rely on the image content, and are computed either in the region inside the contour, or on a region of interest (ROI) that contains the marker. The criteria we used are listed in Table 1.

The computation done on the various contours extracted from the markers allows us to build a data set to be classified. These elements will be denoted as feature vectors whose each dimension corresponds to one feature listed in Table 1.

3.5. Support vector machine

The classification stage is performed by a regular support vector machine (SVM) classifier [15]. Because our problem corresponds to non-linear classification, we used the kernel version of the classifier relying on a Gaussian radial basis function.

While SVM is a commonly used classifier, other well known classifiers like neural networks, decision trees, Bayesian methods, etc. may have been used as well. However, the goal of this paper being the detection of lesions in DBT, our choice went to SVM because it is known to perform well in many applications, including 2D mammography CAD systems [31–33]. The relevance of this choice was confirmed by the experimental results.

The operating points obtained after classification that are presented in this paper are computed using a probability to belong to one class estimated using the approach proposed in [34].

3.6. Results

3.6.1. Validation process

In order to assess the performance of the density detection channel, a cross-validation approach has been implemented.

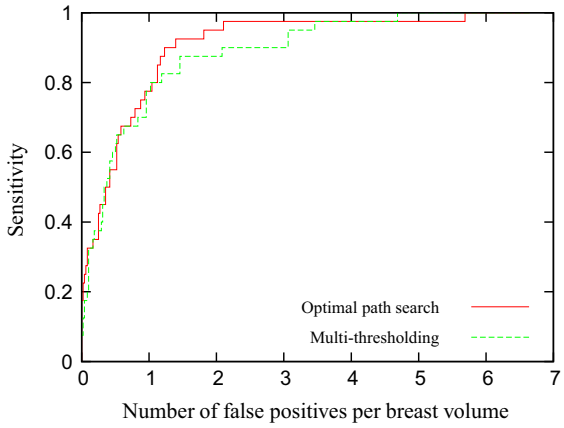

Fig. 8. Performance of the whole dense kernel detection channel.

As illustrated in Fig. 7, for each tomosynthesis volume (x_k), all other volumes ($\{x_j, j \neq k\}$) were used to learn the SVM parameters. The resulting classifier (whose parameters are C, γ) was used to process the remaining volume. The SVM parameters learning stage was done using a leave-one-out scheme. These parameters are the ones that maximize the classifier performance, which is assessed as follows: for all the elements $x_{k'}$ of the learning base ($\{x_j, j \neq k\}$), $x_{k'}$ is processed by the SVM learned with $\{x_j, j \neq k, j \neq k'\}$.

3.6.2. Performance

Fig. 8 presents the performance of the classification stage. Since the operating point of the preliminary marker stage was set to 100%, this figure represents the overall performance of this mass detection channel. The two curves represent the performance associated with the two proposed segmentation methods. Note that the curve obtained using the optimal path search is better than the one using multi-thresholding for high sensitivities. This can be explained by a better quality of the contours provided

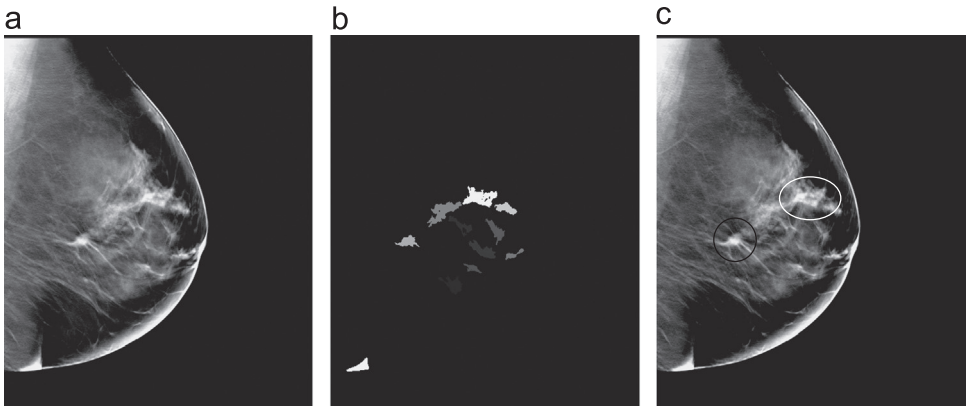

Fig. 9. Example of the full processing performed by the dense kernel detection channel. (a) Plane of a DBT volume that contains a spiculated lesion. (b) Fuzzy detection map obtained after coarse detection. (c) Result after classification: the lesion (circled in black) and a false positive (circled in white) are kept.

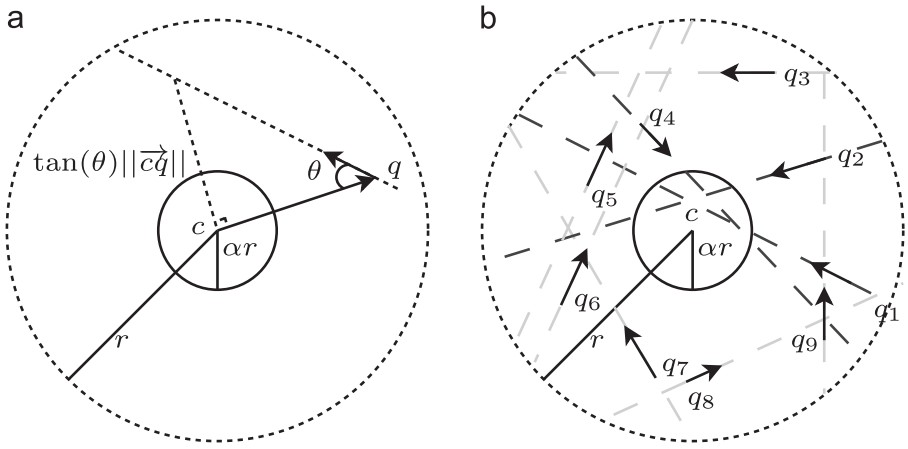

Fig. 10. (a) Nested circles used for the computation of $K_{c,r,q}$ which is here equal to 0. (b) Computation of $Z_{c,r}$: the arrows on the dark dashed lines (passing through q_1, q_2 and q_4) converge while others do not.

by the path-based segmentation method. Actually, the assumption behind the multi-thresholding is that there is no more tissue superimposition in the image. Since digital breast tomosynthesis volumes are built using only a rough approximation of the Radon theorem [35], tissue superimposition is unfortunately not completely removed.

On the same figure, we can see that after segmentation, feature extraction and classification, the mean number of false positives per volume is only 1.23 for a sensitivity of 90% ([76–97] 95% confidence interval) for the optimal path method, which is comparable to state of the art methods [11].

Fig. 9 illustrates the detection channel on a breast containing a spiculated lesion with a dense kernel. In this example, the marker extraction stage provides several suspicious regions, one of them being the lesion. Some of these suspicious zones are discarded during the next stage of the chain. The final output only shows one false positive and the lesion, which appear on the same plane.

4. Detection of lesions without dense kernel

Some of the lesions are not associated with a dense kernel in the image, but rather with a strong convergence pattern. While this kind of findings has partially been addressed in the literature [36,37] for DBT, a full scheme including robust decision making is still needed. In this section, we present the design and the validation of such a channel, which has previously been introduced in Section 3.

4.1. Detection of convergences

In order to detect convergences, we use the *a contrario* approach we previously proposed [36,37] as a marker stage. The idea behind this method is to define a convergence measure and to compute the probability of its values in a naive model, which is supposed to represent the content of a normal breast with no specific convergent structures, and is inspired by the *a contrario* approaches developed e.g. by Desolneux et al. [16,38,39], Muse et al. [40], Cao et al. [41].

In our case, *a contrario* modeling is relevant because it allows setting up an intuitive detection scheme. However, the efficiency of such a method is mainly driven by the quality of the criterion that is used to define the naive model. Therefore *a contrario* reasoning is mainly a tool to interpret this measure.

Our convergence criterion is based on a binary random variable $K_{c,q,r}$ that is equal to 1 if and only if the pixel q, whose underlying structure has an orientation angle θ, "converges" toward the disk of center c and whose radius (αr, $\alpha \in]0;1]$) is proportional to r as defined in Eq. (7) and illustrated in Fig. 10(a)

$$K_{c,q} = \begin{cases} 1 & \text{if } (\alpha r < \|\overrightarrow{cq}\| \leq r) \wedge (\tan(\theta)\|\overrightarrow{cq}\| \leq \alpha r) \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

The convergence measure is then expressed using Eq. (8), which just counts how many pixels converge toward a point c of the image domain Ω as illustrated in Fig. 10(b)

$$Z_{c,r} = \sum_{q \in \Omega} K_{c,q,r} \tag{8}$$

Let us note that while the convergence criterion we use is similar to what has been introduced by Karssemeijer and te Brake [42], it differs in several points. First, we use $\tan(\theta)$ instead of its approximation θ, which enables us to compute convergence for pixels close to the circle of radius αr. Secondly, in our case the convergence criterion is computed for all the pixels, regardless of their signal intensity. This results in a simpler computation which relies on less parameters. A last difference is the computation of the orientations map, as described below.

In our experiments, the naive model assumes that the orientations extracted for each pixel are uniformly and identically distributed. In practice they are computed from non-correlated gradient maps [36]. This model is then used to derive a suitable threshold (λ) for Z. To do that, we consider a parameter ε, which is the maximum mean number of occurrences of the event $Z_{c,r} > \lambda$

$$\sum_{\substack{c \in \Omega \\ r \in [\![R_{\min}; R_{\max}]\!]}} P[Z_{c,r} > \lambda] < \varepsilon \tag{9}$$

with $R_{\min}$ and $R_{\max}$ the minimum and maximum radii to be considered, respectively. They reflect the scale of the targeted structures. Obviously, the parameter ε can be interpreted as the number of false alarms (NFAs) to be accepted in a healthy breast that verifies the naive model.

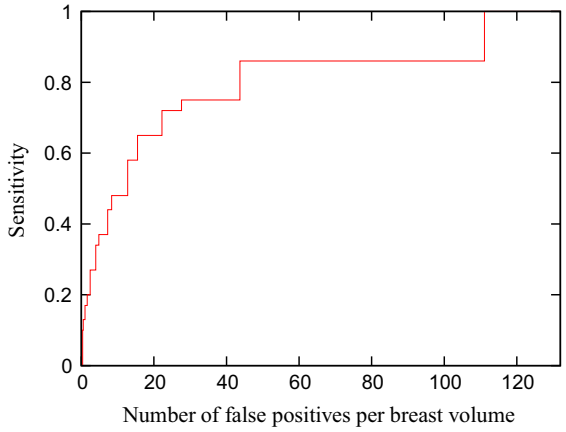

Fig. 11. Performance of the *a contrario* detector for spiculated lesions only.

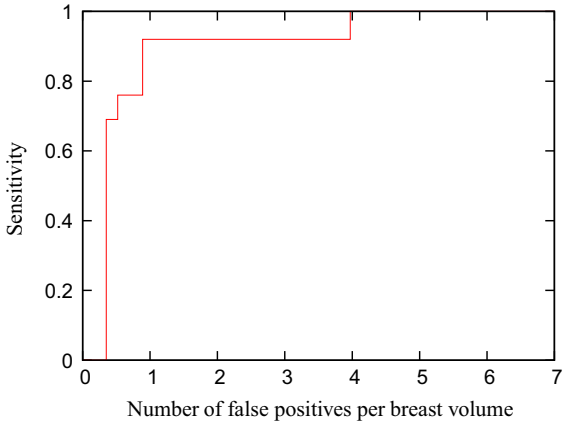

Fig. 12. Performance of the *a contrario* detector for architectural distortions and highly spiculated lesions.

This method allows us to retrieve not only architectural distortions with a stellate pattern but also highly spiculated masses. Because of their morphology, irregular and lobulated lesions are unlikely to be detected by this channel. Furthermore, to keep a high sensitivity with a reasonable specificity, it is not possible to detect all kinds of spiculated lesions with this convergence criterion. This is illustrated in the curve of Fig. 11. Note that in this plot, the scale of the false positives axis has been compressed in order to draw the curve for acceptable sensitivities.

However, if we focus only on stellate lesions, i.e. architectural distortion and highly spiculated lesions whose convergence characteristics are close to architectural distortions (13 lesions in our database), the method becomes suitable to be used as a marker stage as seen in Fig. 12: for instance a sensitivity of 92% is obtained for 0.9 false positives per volume.

The curves of Figs. 11 and 12 were obtained by tuning the parameter ε, which represents the number of false detections per volume we accept to get, assuming the naive model.

4.2. False positive reduction

A classification step similar to the one used for dense kernels detection was implemented to reduce the amount of false positives. The main difference is that all the tomographic planes were used here to make the decision on whether a marker is suspicious or not according to the extracted features. The motivation for this divergence in the analysis is that stellate patterns are rather subtle, hence a better characterization is expected from the analysis of all tomographic planes.

4.2.1. 3D aggregation

In order to take into account the 3D information not used due to the 2D processing, the 2D markers on each plane are gathered together in 3D. To do that, a detection map A is built from the aggregation of all the disks of center c and radius αr for all the pairs (c, r) labeled as suspicious during the first processing stage. Connected components A^j of this 3D map are then labeled. For all j, $\{A_i^j\}$ will denote the set of connected components contained in A^j for each plane intersected by this marker. Each component A^j is then kept if the classification step considers that at least one A_i^j corresponds to a lesion. Furthermore, keeping the whole A^j allows not splitting the lesion if there is a bad classification on an intermediate plane. This aggregation step is illustrated in Fig. 13.

4.2.2. Feature selection

As explained by Palma et al. [36], false positives generated by the *a contrario* modeling correspond to regular fiber crossing. In order to differentiate such structures from potential malignant lesions, features relying on a statistical analysis (entropy) of the orientations lying within the maker neighborhood were used as classifier input. More precisely, the measures were computed in three different regions: in the aggregation of the convergence centers (αr radius disks), on the border of the same zone and inside the aggregation of the convergence regions (disks of radius

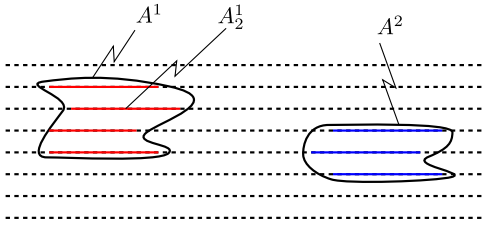

Fig. 13. Illustration of the marker aggregation step for the convergence detection channel.

r). In addition to these measures, a radial gradient index [30] on the border of the convergence region and ratios of the volume of interest dimensions (e.g. $\min(height, width)/depth$) are also computed.

4.3. Performance

The evaluation of the performance of this channel has been conducted similarly to the assessment of the first channel. The 13 lesions presented in Section 2.2 that were used for this purpose prevent us from drawing strong conclusions on the actual performance of this channel. Nonetheless, it allows us to have an idea on its validity.

The performance achieved by this detection scheme is presented in Fig. 14. For high specificity, the sensitivity is better than the one obtained for the first channel. A valid operating point is a sensitivity of 92% ([64–99] 95% confidence interval) for a false positive rate equal to 0.48.

Fig. 15 illustrates an example of the whole convergence detection chain. The image of Fig. 15(a) is a DBT plane containing an architectural distortion. The *a contrario* detection (see Fig. 15(b)) provides four suspicious findings in this plane, one of them being the actual lesion. The false positives reduction stage (see Fig. 15(c)) allows getting rid of two of them. Eventually, two regions are kept: the actual lesion and a false positive.

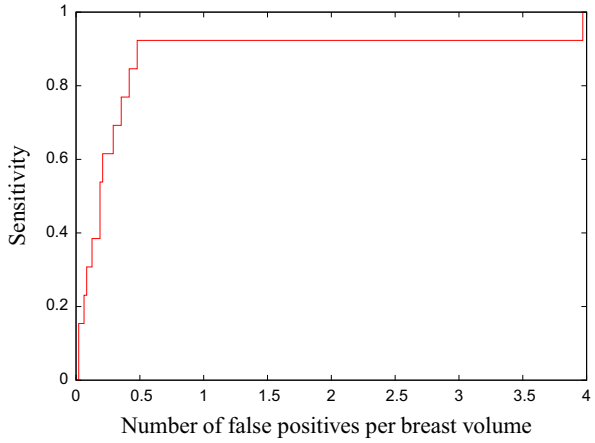

Fig. 14. Performance of the suspicious convergence detection channel.

5. Aggregation and results

In this section, we describe the final aggregation step, detail the experimental results obtained with the proposed complete detection scheme and briefly discuss some implementation details.

5.1. Fusion of the two channels

As explained in the two previous sections, each channel is dedicated to the detection of one type of potential lesion based on its specific characteristics. Therefore the most natural aggregation is a disjunctive fusion, i.e. a region is considered as a potential lesion if it is detected as such in at least one channel. This step can further improve the performance for each class of lesions, since the two considered classes are not perfectly separated and a lesion of one class can actually be detected by the channel dedicated to the other class (this may occur in particular for some spiculated masses).

This approach allows us to compute bounds of sensitivity and specificity by combining the results displayed in Figs. 8 and 14. By merging the two sets of images in the database, for a given specificity, at least a specificity equal to the weighted sum of the two specificity values of the two channels is guaranteed (the weights being the proportion of malignant cases in each set of images).

5.2. Results and discussion

The performance of the complete processing is displayed in Table 2. It has been obtained by using threshold values of the classifier probabilities corresponding to equivalent sensitivity values in the results of Figs. 8 and 14. The results of the two

Table 2
Performances of our 2-channels method and the methods in [11]. Note that the last 3 methods address mass only while our method addresses also architectural distortions and that database was not the same for the evaluation of this last one.

Sensitivity (%)	# of false positives per breast			
	Our method	Chan et al. [11]		
		2D method	3D method	Combined 2D + 3D
80	1.31	2.43	1.46	0.84
90	1.60	3.65	2.52	1.61
96	1.81	>3.65	>3	>2

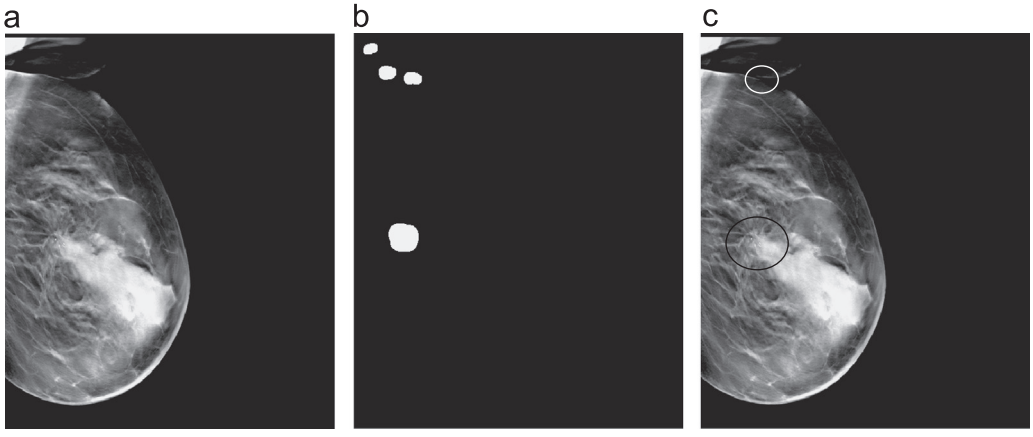

Fig. 15. Example of the full processing of the convergence detection channel. (a) Plane of a DBT volume that contains an architectural distortion. (b) *A contrario* detection result. (c) Result after classification: the lesion (circled in black) and a false positive (circled in white) are kept.

channels are combined by computing the union (logical *or*) of the detection marks obtained in each channel. The validation is performed using a cross-validation method, as done before for each channel, but this time using all cancer and normal cases.

Interestingly enough, the combination of the two channels actually allows us to obtain better performances than the lower bound described before (Section 5.1). For instance this bound would be 1.7 false positives for a sensitivity of 90%, and 1.45 false positives for a sensitivity of 80%. This is explained by the potential detection of strongly spiculated masses (used in the distortion channel evaluation) by the dense kernel channel, as illustrated in Fig. 16. While this lesion was not suspected by the architectural distortion channel, after the false positive reduction step, it was correctly identified by the dense kernel channel.

Unfortunately, there is no publicly available database on which several methods could be evaluated. However, the results obtained by the proposed approach are similar to those published in the literature. Among the works providing performance evaluation, let us mention the one of Chan et al. [5], where a 3D detection method was evaluated on 26 volumes with 13 malignant masses and 10 benign ones, 2 malignant architectural distortions and 1 benign one. The results have shown a sensitivity of 80% (respectively 85%) for 2 (respectively 2.2) false positives per case. However a fair comparison with our results is not possible since the database is not the same, and our approach focuses on malignant cases only.

Chan et al. [43] have then combined this approach with a detection in projection images, with an evaluation on a database containing 41 malignant masses and 11 benign ones. The additional 2D information reduces the number of false positives from 1.6 to 1.19 for a sensitivity of 80%, and from 3.04 to 2.27 for a sensitivity of 90%.

Fig. 16. Plane of a DBT volume exhibiting a strongly spiculated lesion, used in the convergence channel evaluation. Although this lesion is not detected by this channel, it is correctly detected by the dense kernel channel, and therefore by the final fusion step.

More recently, the same group proposed a deeper study of the differences between 2D and 3D detection approaches, and their combination [11], with an evaluation on 100 volumes containing 69 malignant lesions. The comparison with our approach is then more relevant since the databases are of the same order of magnitude (our validation is based on 53 malignant lesions). Let us recall that DBT is still a new acquisition technique resulting in difficulty to collect data. Moreover results are also provided only for the detection of malignant lesions, which corresponds to our configuration. The operating points of the three methods described by Chan et al. [11] are displayed in Table 2. However, let us note that these methods address mass detection only. Nonetheless this table shows direct comparison with our two channels scheme, which also detects architectural distortions to demonstrate that while addressing more findings we may have similar performances and even better ones at high sensitivity.

For 80%, the combined approach of Chan et al. [11] provides a better specificity than our method, while the 2D only and 3D only methods are worse. However, for a 90% sensitivity, our method provides as good results as the combined approach. Moreover, our method allows setting the sensitivity above 90% with a small increase of FP per volume (+13%).

It should be noted that the databases are still quite small in comparison to the high variability of lesions. Furthermore, the databases used by Chan et al. [11] and in this paper are not the same, which makes any comparison difficult. Again, the case of architectural distortions is not addressed by Chan et al. [11]. If we compare the results in Table 2 with those obtained by the first channel only (Fig. 8), it appears that for low sensitivity (80%), the combined method of Chan et al. [11] is slightly better. However, this becomes the opposite when the sensitivity is increased (e.g. for 90%, FP rates are 1.61 vs 1.23). Despite the above-mentioned reasons regarding this comparison, this study shows that our results are comparable to the ones obtained with already published algorithms while addressing more findings.

Let us finally mention three other works. First, Singh [44] has developed a masses detection method based on information theory. Results on a database of 28 lesions (among which 10 are malignant) show a FP rate of 2.4 for a sensitivity of 90%. Secondly, Chan et al. [45] have assessed the influence of the number of projections on the CAD system using a 3D method they previously proposed [5]. This was evaluated on a database of 14 malignant lesions. For a sensitivity of 80%, a FP rate of 0.87 was obtained when using 11 projections, and 1.11 for 21 projections. These results help us to confirm that our system is comparable to what can be found in the literature. Finally, van Schie et al. [3] proposed to train a DBT CAD system with 2D mammograms. While validation was done on a meaningful database, results obtained were slightly worse than our method and the one of Chan et al. [11]. Actually, a sensitivity around 85% (resp. 90%) was obtained at around 2 (resp. 3) false positives per breast volume.

5.3. Implementation details

The prototype used for the evaluation of the proposed methods was implemented in C++. Complexity of the overall method is a combination of every step involved in the approach. For the sake

Table 3
Complexity of the main parts of the proposed CAD scheme.

Step	Complexity	Notes
Fuzzy connected filters	$O(N \log(N))$	N being the number of voxels
Mass segmentation	$O(K)$	K being the number of pixels in the ROI to segment
A contrario detection	$O(NR^2)$	N being the number of voxels and R the diameter of the largest potential lesion

of clarity we summarize the main ones in Table 3. While not fully optimized, the implementation was multi-threaded. Execution time on an 8-core Intel Xeon X5482@3.20 GHz is about 2 min for a 5 cm thick breast volume.

6. Conclusion

In this paper, we have proposed an original method for detecting potential cancerous lesions in DBT images, based on the design of specific methods for masses on one hand, and architectural distortions on the other hand, leading to two detection channels, which are combined in a final decision step. These methods exploit the 3D information provided by the DBT images. Only a few steps are performed on 2D slices, when it is more relevant, and the results are then aggregated in 3D. The results, evaluated on a database of 101 images, containing 56 malignant lesions, are very promising and compare favorably to those published in the literature, especially for high sensitivity operating points. They constitute a proof of concept, and the proposed method can be the basis for a complete CAD system dedicated to DBT volumes.

Further improvements could be proposed. The imprecision induced by the DBT artefacts could be modeled using the notion of fuzzy umbra images used in the first channel. While the mathematical and algorithmic aspects can be dealt with, an increase in computation time would however be expected. Imprecision and uncertainty could also be modeled in the segmentation and classification steps.

As for the second channel, it would be interesting to investigate a more realistic modeling of the tomographic planes content, for instance by taking into account the correlation between voxels.

Further work aims at an extended evaluation on a larger database, which would also enable to include more lesion types in the proposed CAD system.

Conflict of interest

None declared.

Acknowledgments

This work was partially funded by a grant from ANRT (1165/2006) during the Ph.D. thesis of G. Palma at Telecom ParisTech and GE Healthcare.

References

[1] T. Wu, A. Stewart, M. Stanton, T. McCauley, W. Phillips, D.B. Kopans, R.H. Moore, J.W. Eberhard, B. Opsahl-Ong, L. Niklason, M.B. Williams, Tomographic mammography using a limited number of low-dose cone-beam projection images, Med. Phys. 30 (2003) 365–380.
[2] S. Bernard, S. Muller, J. Onativia, Computer-aided microcalcification detection on digital breast tomosynthesis data: a preliminary evaluation, in: International Workshop on Digital Mammography (IWDM), 2008, pp. 151–157.
[3] G. van Schie, M.G. Wallis, K. Leifland, M. Danielsson, N. Karssemeijer, Mass detection in reconstructed digital breast tomosynthesis volumes with a computer-aided detection system trained on 2D mammograms, Med. Phys. 40 (2013) 041902.
[4] H.-P. Chan, J. Wei, B. Sahiner, E.A. Rafferty, T. Wu, M.A. Roubidoux, R.H. Moore, D.B. Kopans, L.M. Hadjiiski, M.A. Helvie, Computerized detection of masses on digital tomosynthesis mammograms—a preliminary study, in: International Workshop on Digital Mammography (IWDM), 2004, pp. 199–202.
[5] H.-P. Chan, J. Wei, B. Sahiner, E.A. Rafferty, T. Wu, M.A. Roubidoux, R.H. Moore, D.B. Kopans, L.M. Hadjiiski, M.A. Helvie, Computer-aided detection system for breast masses on digital tomosynthesis mammograms: preliminary experience, Radiology 237 (2005) 1075–1080.
[6] I. Reiser, R.M. Nishikawa, M.L. Giger, T. Wu, E.A. Rafferty, R. Moore, D.B. Kopans, Computerized mass detection for digital breast tomosynthesis directly from the projection images, Med. Phys. 33 (2006) 482–491.
[7] I. Reiser, R. Nishikawa, M. Giger, D. Kopans, E. Rafferty, T. Wu, R. Moore, A multi-scale 3D radial gradient filter for computerized mass detection in digital tomosynthesis breast images, in: Computer Assisted Radiology and Surgery (CARS), vol. 1281, Berlin, Germany, 2005, pp. 1058–1062.
[8] G. Peters, Computer-aided detection for digital breast tomosynthesis (Ph.D. thesis), Ecole Nationale Supérieure des Télécommunications, 2007.
[9] G. Peters, S. Muller, B. Grosjean, S. Bernard, I. Bloch, A hybrid active contour model for mass detection in digital breast tomosynthesis, in: SPIE Symposium on Medical Imaging, vol. 6514-1V, 2007, pp. 1–11.
[10] G. Palma, G. Peters, S. Muller, I. Bloch, Masses classification using fuzzy active contours and fuzzy decision trees, in: SPIE Symposium on Medical Imaging, vol. 6915, San Diego, CA, USA, 2008, pp. 1–11.
[11] H.-P. Chan, J. Wei, Y. Zhang, M.A. Helvie, R.H. Moore, B. Sahiner, L. Hadjiiski, D.B. Kopans, Computer-aided detection of masses in digital tomosynthesis mammography: comparison of three approaches, Med. Phys. 35 (2008) 4087–4095.
[12] S. Singh, G.D. Tourassi, J.A. Baker, E. Samei, J.Y. Lo, Automated breast mass detection in 3D reconstructed tomosynthesis volumes: a featureless approach, Med. Phys. 35 (2008) 3626–3636.
[13] S. Singh, G.D. Tourassi, A.S. Chawla, R.S. Saunders, E. Samei, J.Y. Lo, Computer-aided detection of breast masses in tomosynthesis reconstructed volumes using information-theoretic similarity measures, in: SPIE Symposium on Medical Imaging, vol. 6915, 2008, pp. 1–8.
[14] G. Palma, I. Bloch, S. Muller, Fast fuzzy connected filter implementation using max-tree updates, Fuzzy Sets Syst. 161 (2010) 118–146.
[15] N. Cristianini, J. Shawe-Taylor, An Introduction to Support Vector Machines and Other Kernel-based Learning Methods, Cambridge University Press, New York, NY, USA, 2000.
[16] A. Desolneux, L. Moisan, J. Morel, Meaningful alignments, Int. J. Comput. Vis. 40 (2000) 7–23.
[17] A.H. Andersen, A.C. Kak, Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm, Ultrason. Imaging 6 (1984) 81–94.
[18] G. Palma, I. Bloch, S. Muller, Fuzzy connected filters for fuzzy gray scale images, in: International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU), Malaga, Spain, 2008, pp. 667–674.
[19] A. Rosenfeld, Fuzzy digital topology, Inf. Control 40 (1979) 76–87.
[20] O. Nempont, J. Atif, E. Angelini, I. Bloch, A new fuzzy connectivity class. Application to structural recognition in images, in: Discrete Geometry for Computer Imagery, Lyon, 2008, pp. 446–557.
[21] O. Nempont, J. Atif, E. Angelini, I. Bloch, A new fuzzy connectivity measure for fuzzy sets and associated fuzzy attribute openings, J. Math. Imaging Vis. 34 (2009) 107–136.
[22] U. Braga-Neto, J. Goutsias, A theoretical tour of connectivity in image processing and analysis, J. Math. Imaging Vis. 19 (2003) 5–31.
[23] A. Rosenfeld, The fuzzy geometry of image subsets, Pattern Recognit. Lett. 2 (1984) 311–317.
[24] I. Bloch, H. Maitre, Fuzzy mathematical morphologies: a comparative study, Pattern Recognit. 28 (1995) 1341–1387.
[25] G. Palma, I. Bloch, S. Muller, R. Iordache, Fuzzifying images using fuzzy wavelet denoising, in: FUZZ-IEEE, Jeju, Korea, 2009, pp. 135–140.
[26] J. Beutel, H.L. Kundel, R.L. Van Metter, Handbook of Medical Imaging. Physics and Psychophysics, vol. 1, SPIE Press, Washington, USA, 2000.
[27] ACR, BI-RADS Breast Imaging Reporting and Data System, American College of Radiology, 2003.
[28] S. Timp, N. Karssemeijer, A new 2D segmentation method based on dynamic programming applied to computer aided detection in mammography, Med. Phys. 31 (2004) 958–971.
[29] L. Apffel, G. Palma, I. Bloch, S. Muller, Fuzzy segmentation of masses in digital breast tomosynthesis images based on dynamic programming, in: International Conference on Imaging Theory and Applications (IMAGAPP), Angers, France, 2010, pp. 7–13.
[30] M.A. Kupinski, M.L. Giger, Automated seeded lesion segmentation on digital mammograms, IEEE Trans. Med. Imaging 17 (1998) 510–517.
[31] A. Cao, Q. Song, X. Yang, S. Liu, C. Guo, Mammographic mass detection by vicinal support vector machine, in: IEEE International Joint Conference on Neural Networks, vol. 3, 2004, pp. 1953–1958.
[32] R. Campanini, D. Dongiovanni, E. Iampieri, N. Lanconelli, M. Masotti, G. Palermo, A. Riccardi, M. Roffilli, A novel featureless approach to mass detection in digital mammograms based on support vector machines, Phys. Med. Biol. 49 (2004) 961–975.
[33] H. Bornefalk, Use of quadrature filters for detection of stellate lesions in mammograms, in: Scandinavian Conference on Image Analysis (SCIA), 2005, pp. 649–658.
[34] T. Wu, C. Lin, R. Weng, Probability estimates for multi-class classification by pairwise coupling, J. Mach. Learn. Res. 5 (2004) 975–1005.
[35] J. Radon, Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten, Ber. Ver. Sachs. Akad. Wiss. Leipz., Math. Phys. Kl. 69 (1917) 262–277.
[36] G. Palma, S. Muller, I. Bloch, R. Iordache, Convergence areas detection in digital breast tomosynthesis volumes using a contrario modeling, in: SPIE Symposium on Medical Imaging, Lake Buena Vista, FL, USA, 2009, pp. 1–8.

[37] G. Palma, S. Muller, I. Bloch, R. Iordache, Fast detection of convergence areas in digital breast tomosynthesis, in: IEEE International Symposium on Biomedical Imaging (ISBI), Boston, MA, USA, 2009, pp. 847–850.
[38] A. Desolneux, L. Moisan, J.-M. Morel, Edge detection by Helmholtz principle, J. Math. Imaging Vis. 14 (2001) 271–284.
[39] A. Desolneux, L. Moisan, J. Morel, A grouping principle and four applications, IEEE Trans. Pattern Anal. Mach. Intell. (PAMI) 25 (2003) 508–513.
[40] P. Muse, F. Sur, F. Cao, Y. Gousseau, J.-M. Morel, An a contrario decision method for shape element recognition, Int. J. Comput. Vis. 69 (2006) 295–315.
[41] F. Cao, J. Delon, A. Desolneux, P. Muse, F. Sur, A unified framework for detecting groups and application to shape recognition, J. Math. Imaging Vis. 27 (2007) 91–200.
[42] N. Karssemeijer, G.M. te Brake, Detection of stellate distortions in mammograms, IEEE Trans. Med. Imaging 5 (1996) 611–619.
[43] H.-P. Chan, J. Wei, Y. Zhang, R.H. Moore, D.B. Kopans, L. Hadjiiski, B. Sahiner, M.A. Roubidoux, M.A. Helvie, Computer-aided detection of masses in digital tomosynthesis mammography: combination of 3D and 2D detection information, in: SPIE Symposium on Medical Imaging, vol. 6514, 2007, pp. 1–6.
[44] S. Singh, Computer aided detection of masses in breast tomosynthesis imaging using information theory principles (Ph.D. thesis), Duke University, 2008.
[45] H.-P. Chan, J. Wei, Y. Zhang, B. Sahiner, L. Hadjiiski, M.A. Helvie, Detection of masses in digital breast tomosynthesis mammography: effects of the number of projection views and dose, in: International Workshop on Digital Mammography (IWDM), Springer-Verlag, Berlin, Heidelberg, 2008, pp. 279–285.

Giovanni Palma received the engineering degree in computer science from EPITA, France, in 2005, the master and Ph.D. degrees in signal and image processing from Telecom ParisTech, France, in 2006 and 2010, respectively. His research interests include image processing, mammography, computer aided detection and digital breast tomosynthesis.

Isabelle Bloch is graduated from the Ecole des Mines de Paris, Paris, France, in 1986, she received the Master's degree from the University Paris 12, Paris, in 1987, the Ph.D. degree from the Ecole Nationale Supérieure des Télécommunications (Telecom ParisTech), Paris, in 1990, and the Habilitation degree from the University Paris 5, Paris, in 1995.

She is currently a Professor with the Signal and Image Processing Department, Telecom ParisTech, in charge of the Image Processing and Understanding Group. Her research interests include 3D image and object processing, computer vision, 3D and fuzzy mathematical morphology, information fusion, fuzzy set theory, structural, graph-based, and knowledge-based object recognition, spatial reasoning, and medical imaging.

Serge Muller is the chief engineer of breast care business at GE Healthcare. He was graduated from Paris University in 1986 with a Ph.D. degree in spatial technologies and astronomy, and in 2004 in applied mathematics. His current domains of interest are image processing, spectral mammography, digital breast tomosynthesis and robotics.