

Fast and effective quantification of symmetry in medical images for pathology detection: Application to chest radiography

Laurens Hogeweg, Clara I. Sánchez, Pragnya Maduskar, Rick H.H.M. Philipsen, and Bram van Ginneken^{a)}

Diagnostic Image Analysis Group, Radboud University Medical Center, Nijmegen 6521GA, The Netherlands

(Received 20 November 2016; revised 13 January 2017; accepted for publication 15 January 2017; published 4 May 2017)

Purpose: Symmetry is an important feature of human anatomy and the absence of symmetry in medical images can indicate the presence of pathology. Quantification of image symmetry can then be used to improve the automatic analysis of medical images.

Methods: A method is presented that computes both local and global symmetry in 2D medical images. A symmetry axis is determined to define for each position p in the image a mirrored position p' on the contralateral side of the axis. In the neighborhood of p' , an optimally corresponding position p_s is determined by minimizing a cost function d that combines intensity differences in a patch around p and the mirrored patch around p_s and the spatial distance between p' and p_s . The optimal value of d is used as a measure of local symmetry s . The average of all values of s , indicated as S , quantifies global symmetry. Starting from an initial approximation of the symmetry axis, the optimal orientation and position of the axis is determined by greedy minimization of S .

Results: The method was evaluated in three experiments concerning abnormality detection in frontal chest radiographs. In the first experiment, global symmetry S was used to discriminate between 174 normal images and 174 images containing diffuse textural abnormalities from the publicly available CRASS database of tuberculosis suspects. Performance, measured as area under the receiver operating characteristic curve A_z was 0.838. The second experiment investigated whether adding the local symmetry s as an additional feature to a set of 106 texture features resulted in improvements in classifying local patches in the same image database. We found that A_z increased from 0.878 to 0.891 ($P = 0.001$). In the third experiment, it was shown that the contrast of pulmonary nodules, obtained from the publicly available JSRT database, increased significantly in the local symmetry map compared to the original image.

Conclusions: We conclude that the proposed algorithm for symmetry computation provides informative features which can be used to improve abnormality detection in medical images both at a local and a global level. © 2017 American Association of Physicists in Medicine [https://doi.org/10.1002/mp.12127]

Key words: chest radiography, computer-aided detection, symmetry, symmetry axis

1. INTRODUCTION

Symmetry, an ubiquitous property of both natural and man-made objects, is the property of an object being invariant to certain types of transformations. The most well-known forms of symmetry are reflection, or bilateral symmetry, and rotation symmetry. Symmetry as a general feature of objects has been extensively studied in computer vision. Being such a fundamental property of many objects, there have been numerous applications where symmetry has been applied, for example in face detection,¹ object tracking,^{2,3} and analysis of textures.⁴ A detailed overview of many aspects of symmetry computation and its applications can be found elsewhere.⁵ Many of the proposed algorithms considered symmetrical properties in objects described by their boundaries.^{6–8} More recently techniques have been developed that detect symmetry directly in images, using point descriptors (features) to measure similarity between symmetric points.^{9,10}

Although the output of these methods is mostly a binary measure, a continuous symmetry measure contains more

information as it can be used to impose an ordering on a series of objects such that an object with a smaller measure is judged to be less symmetric. A well-known example of such a measure is the continuous symmetry measure (CSM) by Zabrodsky *et al.*,⁸ which quantifies symmetry in object boundaries. A number of papers have used continuous symmetry measures, such as the CSM, to find correlations with other characteristics. Examples are the relation between facial symmetry and subjective measures of attractiveness,¹¹ fluctuating asymmetry and developmental instability,¹² and molecule symmetry and enzyme activity.¹³

The human body exhibits a large degree of symmetry, clearly visible on the outside, and numerous organs such as the brain, lungs, and visual system also display symmetry. A loss of symmetry in these organs is often an indication of a disturbance of their normal functioning. For this reason, visual assessment of symmetry in medical images is typically used by human specialists for image interpretation and pathology detection. In automated medical image analysis, the use of symmetry has been limited and mainly focused on

brain MRI. In Liu *et al.*,¹⁴ symmetry was used to robustly extract the midsagittal plane in pathological brain images. In Sun *et al.*,^{15,16} the detection, segmentation, and classification of brain lesions was performed using a symmetry measure that involves computing point-to-point similarities based on the curvature of the gradient vector flow. Digital subtraction techniques have been also used to show differences between the two sides of the symmetry axis. Li *et al.*^{17,18} performed registration of the left and right lung fields in posterior–anterior chest radiographs, followed by subtraction, to suppress normal symmetrical structures and enhance pathology. In a later study,¹⁹ a similar technique was successfully used as a postprocessing stage to reduce the number of false-positive detections in a CAD scheme to detect nodules.

However, many previously proposed methods^{9,15,16} are formulated from an assumption of perfect symmetry, but are applied to images that have varying degrees of symmetry. Many natural objects clearly display properties of symmetry, but this symmetry is usually not perfect.²⁰ In the human body, symmetrical organs, such as the brain, are not perfectly symmetrical, even in healthy subjects.²¹ The lungs are not fully symmetrical (e.g., the left lung has two lobes and the right one three and the shadow of the heart breaks the symmetry in chest radiographs), but still exhibit a large amount of symmetry in how its internal structures, such as vessels and airways, are organized. The amount of symmetry in a medical image also depends on the properties of the imaging device. Projection radiography, computed tomography (CT), and magnetic resonance imaging (MRI) have different resolutions, contrast, etc. A method analyzing symmetry in medical images should be able to deal with this inherent normal asymmetry.

In this paper, we propose a new algorithm to assess symmetry in 2D medical images with the aim to detect the presence of pathology. The innovative contribution of the proposed method is twofold. First, unlike previous methods, the proposed algorithm deals with both the inherent normal asymmetry of the organs and asymmetry as a result of pathology. The normal asymmetry cannot be corrected for using existing fixed point-to-point analysis or registration methods. Therefore, the method uses point descriptors and similarity measures to describe the image contents and its symmetry.^{22–24} Second, two symmetry measures are provided: a local symmetry measure, and a global symmetry measure derived from local symmetry, which will allow not only determining the presence of pathology but also provide an exact localization of the lesion. These innovative aspects lead to quantification of symmetry perception in medical images without being wrongly influenced by anatomical asymmetries. The potential usefulness of the algorithm is demonstrated in a number of applications involving the detection of abnormalities in chest radiographs where the proposed symmetry measures are correlated to clinical characteristics.

The paper is organized as follows. Section 2 describes the method and its implementation. Section 3 describes the experiments and shows results for three different tasks. Section 4 discusses the results and in Section 5 we conclude.

2. MATERIALS AND METHODS

In this section, the algorithm to obtain a continuous local symmetry and global symmetry measure is described. Local symmetry s is defined as the minimum dissimilarity between two corresponding points on both sides of the symmetry axis. This is a different definition than the one used in some previous works,⁹ where local symmetry is described as the presence of symmetrical structures in a subregion of the full image. On the other hand, global symmetry S is defined as the degree of symmetry in the whole image.

2.A. Prerequisites

The method operates on a discretized 2D gray value image I in which one mirror symmetric object or region of interest is present. Posterior–anterior chest radiographs and brain MRI or CT sections are common medical examples of such images.

Let l be the vertical symmetry axis of the object in I with location $x = x_s$. The image is then divided by l into two sets of points: on the left \mathbf{P}_L and on the right \mathbf{P}_R (Fig. 1(a)). An equal number of points in both sets is not required and the correspondences between their points are not known. \mathbf{P}_L and \mathbf{P}_R can cover all the points in the image I (i.e., $\mathbf{P}_L \cup \mathbf{P}_R = I$) or be subsets of points ($\mathbf{P}_L \subset I$ and $\mathbf{P}_R \subset I$, $\mathbf{P}_L \cup \mathbf{P}_R \neq I$), for example, presegmented structures of interest such as the lung fields in chest radiographs (Fig. 1(a)).

In order to deal with objects which are known *a priori* to exhibit only partial symmetry, we take into account points $p \in \mathbf{P}_L$ (similarly $p \in \mathbf{P}_R$) if $p' \in \mathbf{P}_R$ (similarly $p' \in \mathbf{P}_L$), where p' is the corresponding point in the mirrored position of p with respect to l (see Fig. 1(b)). Given a set $\mathbf{P} \subset \mathbf{P}_L$ (similarly $\mathbf{P} \subset \mathbf{P}_R$), let $\mathbf{P}' \subset \mathbf{P}_R$ (similarly $\mathbf{P}' \subset \mathbf{P}_L$) be the corresponding mirrored set of \mathbf{P} with respect to l (see Fig. 1(a)). The regions formed by \mathbf{P} and \mathbf{P}' are mirror images of each other and are used to compute local symmetry.

2.B. Local symmetry

We define the local symmetry $s(p;l)$ of point $p \in \mathbf{P}$ with coordinates (x,y) given the symmetry axis l as the minimum dissimilarity between p and the points on the contralateral side of l :

$$s(p;l) = \min_{q \in \mathbf{P}'} d(p,q) \quad (1)$$

where $d(p,q)$ is the dissimilarity between points p and q . Higher values of s indicate less similarity between points.

Note that the algorithm considers all points on the other side of the symmetry axis to compute $s(p;l)$, instead of using only the mirrored point p' of p with respect to l as in Fig. 1(a). In this way, the inherent asymmetry in the image is taken into account, in contrast to previously described methods.^{9,15}

Let $\mathbf{f}(p)$ be a point descriptor which describes the local properties of p and consists of an image component $\mathbf{f}_{\text{img}}(p)$ and a position component $\mathbf{f}_{\text{pos}}(p)$:

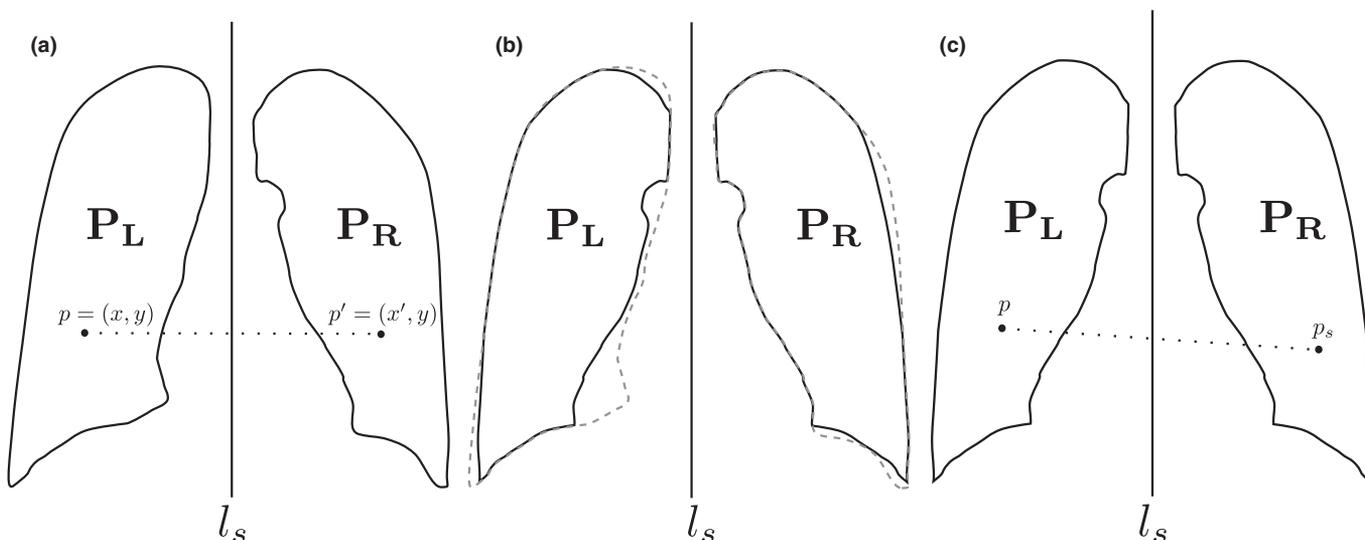


FIG. 1. Computation of local symmetry $s(p)$ for a point p . (a) Point p in \mathbf{P}_L with mirror symmetric point p' in \mathbf{P}_R . (b) Computation is only performed in mirror symmetric sets of locations. The continuous line delimits points \mathbf{P} with corresponding mirror symmetric points \mathbf{P}' . The dashed line represents the initial regions \mathbf{P}_L and \mathbf{P}_R . (c) The optimal matching point p_s is determined by minimizing a cost function d that measures the dissimilarity of position and image characteristics of p with all points in \mathbf{P}_R . The minimum value of d determines $s(p)$. Note that p_s is not necessarily equal to p' .

$$\mathbf{f}(p) = \begin{pmatrix} \mathbf{f}_{\text{img}}(p) \\ \mathbf{f}_{\text{pos}}(p) \end{pmatrix}. \tag{2}$$

The image component \mathbf{f}_{img} can be provided by any suitable local descriptor. Examples are SIFT,²³ local binary patterns,²⁵ etc. In this paper, we define the elements of \mathbf{f}_{img} as the pixel intensities sampled from a square patch Ω around p with patch size (edge length) m pixels, as in Avni et al.,²² and a stride of 1 pixel. As m determines in physical space the area of the regions that are compared, the optimal value depends on the application. From the algorithmic side, m^2 is also the number of elements of the image component of the point descriptor and as such influences the speed of the algorithm (see below). Therefore, for selected applications, a stride larger than 1 pixels might be necessary, but in the experiments in the paper we found the default stride of 1 pixel to be sufficient. In order to determine similar looking positions on both sides of the symmetry axis, computation of $\mathbf{f}_{\text{img}}(q)$ for $q \in \mathbf{P}'$ are calculated using a mirrored version of Ω . In square patches, and with a row-based sampling, mirroring is performed by reversing the order of elements in \mathbf{f}_{img} . The optimal value of m depends on the application and is determined in Section 3.

The position component \mathbf{f}_{pos} consists of the position of p with respect to the symmetry axis l :

$$\mathbf{f}_{\text{pos}}(p) = \begin{pmatrix} x_l \\ y_l \end{pmatrix}, \tag{3}$$

where

$$x_l = \begin{cases} x & \text{if } x \leq x_s \\ 2x_s - x & \text{if } x > x_s \end{cases} \tag{4}$$

and $y_l = y$. The inclusion of the position component in the descriptor ensures that two points are only considered similar

if they have similar visual characteristics and they are approximately spatially symmetrical.

The dissimilarity $d(p, q)$ between two points $p \in \mathbf{P}$ and $q \in \mathbf{P}'$ is then defined as the distance K between their descriptors:

$$d(p, q) = K(\mathbf{f}(p), \mathbf{f}(q)) \tag{5}$$

Many distance definitions are available for K . In this study, the Euclidean distance $K = \|\mathbf{f}_1 - \mathbf{f}_2\|$ is used.

The elements of the point descriptor contain intensity as well as position values. These values were normalized to zero mean and unit standard deviation over both lung fields before computing distances. In order to control the relative contribution of image and position components, a position weight factor w is introduced:

$$\mathbf{f}(p) = \begin{pmatrix} \mathbf{f}_{\text{img}}(p) \\ w\mathbf{f}_{\text{pos}}(p) \end{pmatrix}. \tag{6}$$

High values of w favor matching points which have similar symmetric locations. The value of w depends on the application and optimal values are studied in Section 3.

Determining the local symmetry using Eq. (1) can be computationally expensive if the number of points is large. To reduce computational requirements, finding the most similar point is formulated as a 1-nearest-neighbor problem. Efficient solutions to this problem have been developed which precompute data structures and provide a fast approximation close to the exact solution.²⁶ In this work, we use the approximate nearest neighbor (ANN) algorithm described in Arya and Mount.²⁷ The algorithm uses pre-computed kd-trees and provides an approximate solution which insures that the distance to the approximated nearest neighbor is smaller than $(1+\epsilon)$ times the distance to the true neighbor. For asymmetry computation $\epsilon = 2.0$ is used. In previous work, we established that this value showed

no change in classification performance compared to $\epsilon = 0$ (Ref.²⁸)

2.C. Global symmetry

The global symmetry measure $S(I;l)$ of an image I given the symmetry axis l and the sets \mathbf{P} and \mathbf{P}' is computed by averaging all the local symmetry measures $s(p;l)$ in \mathbf{P} and \mathbf{P}' . If the set of all N locations on both sides of the symmetry axis is defined as $\mathcal{P} = \mathbf{P} \cup \mathbf{P}'$, S is then defined as:

$$S(I;l) = \frac{1}{N} \sum_{p \in \mathcal{P}} s(p;l) \tag{7}$$

Low values of S indicate overall similarity of image characteristics on both sides of the symmetry axis. A value of $S = 0$ indicates that for every point a perfect analog has been found on the other side at the expected reflected position. High values of S indicate the presence of image characteristics on one side which cannot be found on the other side.

Local symmetry values s are spatially correlated and contain redundant information. The global symmetry can therefore be estimated using only a subset of \mathcal{P} without losing its discriminative properties. In Section 3, we determine the effect of reducing the number of samples in \mathcal{P} by introducing a subsample factor k as parameter, which was implemented by sampling every \sqrt{k} th pixel in the x - and y -direction. Pilot experiments showed negligible influence of κ on the optimal values of m or w ; increasing the value of κ is mainly of interest for reduction of computation times.

2.D. Determination of optimal symmetry axis

In medical imaging, the scanning protocol typically ensures that anatomical structures have a fixed orientation and location in the image. For example, in posterior–anterior chest radiographs, the lung fields are centered and the caudocranial direction of the patient is aligned with the y

dimension of the image. For brain imaging with computed tomography or magnetic resonance imaging, a similar fixed relation between patient and image coordinate systems is common. In practice, locations of axes or planes of symmetry are not exactly known and are not necessarily aligned with the image axes.

We estimate the optimal position of the symmetry axis l from an initial approximation by minimization of the global symmetry value S . Note that the ensuing discussion relates to 2D images, but the procedure can be easily extended to higher dimensions.

2.D.1. In-plane rotation

If the patient is rotated in the xy plane, it will cause the symmetry axis to deviate from the verticality which is expected by the algorithm. In order to identify the optimal angle of the symmetry axis, anatomical structures are rotated upright by artificially imposing a range of rotations with different angles to the image. Rotation was performed around the center of the image. The angle which results in the minimum global symmetry S corresponds to the optimal rotation of the image and the upright position of the thorax. Let α be the angle used to rotate the image I in the xy plane around the image center. Let $S(I;l,\alpha)$ denote the global symmetry value computed for I after rotation. The optimal angle α_{opt} is defined as:

$$\alpha_{opt} = \underset{\alpha \in A}{\operatorname{argmin}} S(I;l,\alpha) \tag{8}$$

where A is a set of test angles. Figure 2 shows an example of in-plane rotation for a chest radiograph.

2.D.2. Symmetry axis x -coordinate

The x -position of the initial approximation of the symmetry axis x_s (Section 5) can be displaced from its optimal position. A similar minimization procedure as for the in-plane rotation was used to find the optimal position. Let $S(I;l,\delta)$

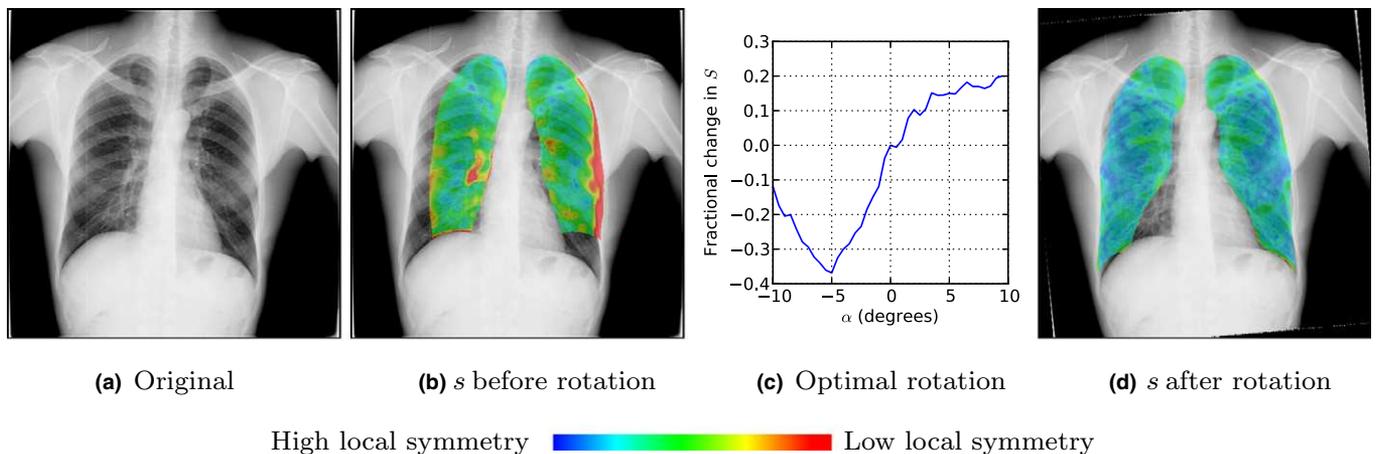


FIG. 2. Example of local symmetry (s) maps before and after optimal rotation in the xy plane. Overlay maps in (b) and (d) indicate local symmetry. Overlay scale has arbitrary units, because s values have only a relative interpretation. The graph in (c) indicates the relation between the rotation angle α and global symmetry S (compared to the original image). The image rotated by the angle corresponding to the minimum value of S in (c) is used in (d). [Color figure can be viewed at wileyonlinelibrary.com]

denote the global symmetry value of image I after applying an horizontal displacement δ to the symmetry axis l . The optimal horizontal displacement δ_{opt} of l is computed as

$$\delta_{opt} = \underset{\delta \in \Delta}{\operatorname{argmin}} S(I; l, \delta) \tag{9}$$

where Δ is a set of test locations.

2.E. Symmetry computation in chest radiographs

In this paper, we select the analysis of chest radiographs to evaluate the performance of the proposed symmetry measures in real medical images. Specific details for symmetry computation in chest radiographs, which are used in the experiments, are given in this section.

The expected scale of normal and abnormal structures determine the working resolution of the images for symmetry computation and the scale at which f_{img} is computed. Chest radiographs were resampled to a fixed width of 512 pixels in all experiments. Optionally, images can be preprocessed to enhance certain structures. The use of a normalization procedure and its properties may influence the symmetry measures. In Section 2.F.3, we evaluate the use of a local normalization procedure which enhances contrast, removes low frequency variations, and normalizes edge strengths.²⁹ This procedure locally normalizes the intensity deviation from the average to the local standard deviation:

$$I_{LN} = \frac{I - \tilde{I}}{\sqrt{\tilde{I}^2 - (\tilde{I})^2}} \tag{10}$$

where I indicates the original image, I_{LN} the locally normalized image and $(\tilde{\cdot})$ blurring by convolution with a Gaussian kernel with scale σ_{LN} . The parameter σ_{LN} is application dependent. The procedure is implemented by applying Eq. (7) to each point in the image.

An initial location of the vertical symmetry axis is determined as follows. The existence of a binary segmentation of the lung fields is assumed, where lung fields have value 1 otherwise 0 (Fig. 3a). A one-dimensional projection image $\bar{y}(x)$ is created by orthogonal averaging over the y direction (Fig. 3b). The function \bar{y} contains two peaks, corresponding to the lung fields, and three valleys corresponding to the two parts at the sides of the image and the part between the lung fields. The approximate position of the symmetry axis is determined by the minimum value in the middle 20% of the curve (gray area in Fig. 3b) or, if there is no single minimum, by the midpoint of the zero region.

Computation of the symmetry measures was performed using locations in the symmetric lung fields only (as in Fig. 1(b)). Please note that symmetric locations P and P' , which depend on l , are redetermined in each iteration of the symmetry axis optimization, including a recomputation of the x -location of the symmetry axis using the procedure described in the previous paragraph during rotation optimization. Patches close to the lung border, with some of the pixels in the patch outside \mathcal{P} , are included in the computation. In chest

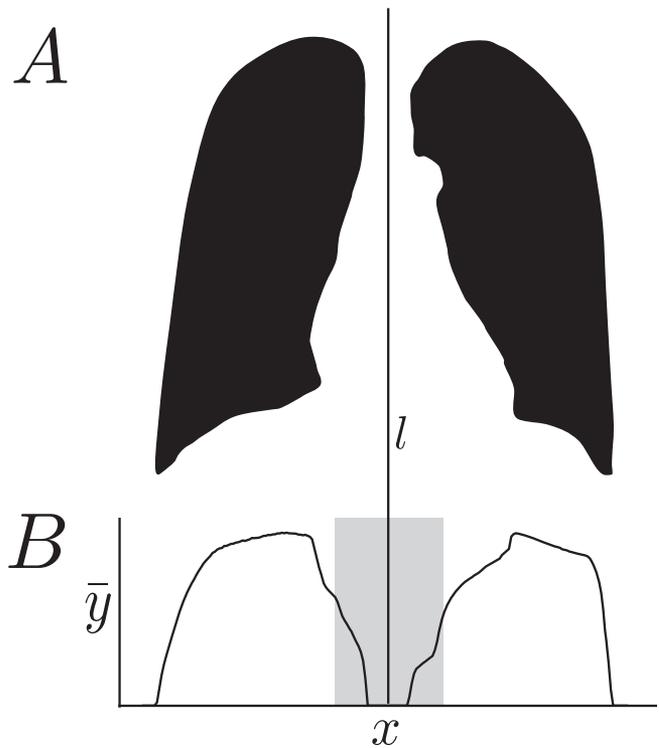


FIG. 3. Computation of initial symmetry axis location. (a) Binary lung segmentation (black = 1, white = 0) (b) Orthogonal projection over y -axis $\bar{y}(x)$. The x -coordinate of the symmetry axis is derived from the gray area.

radiographs this does not introduce artifacts because the symmetry of the image continues outside, but close to, the lung border.

2.F. Experiments

We performed three experiments to validate the proposed method for the detection of pathologies in chest radiographs using global and local symmetry measures. In the first experiment, global symmetry was used to discriminate between normal and abnormal images. In this experiment, the influence of algorithm parameters on the final result was extensively studied. In the second experiment, the contribution of the local symmetry measure to a set of general texture features was determined in an image patch classification task. In the third experiment, the local symmetry measure was used to enhance nodule contrast.

2.F.1. Global symmetry to discriminate between normal and abnormal images

The global symmetry measure S quantifies the presence of symmetrical structures on both sides of the symmetry axis. Chest radiographs (CXRs) of healthy persons are largely symmetric and tend to give low values. The presence of abnormalities in the lung fields tends to increase S . In this experiment, we study the discriminatory power of S to distinguish between normal and abnormal images.

A set of 348 CXRs (174 normal, 174 containing textural abnormalities) was selected from a database consisting of images of tuberculosis (TB) suspects. Images from digital chest radiography units were used (Delft Imaging Systems, the Netherlands) of varying resolutions, with a typical resolution of 1800×2000 pixels, the pixel size was $250 \mu\text{m}$ isotropic. The set is a subset of a publicly available database described in Hogeweg *et al.*³⁰ where normal images and images containing textural abnormalities inside the lung field were selected. The normal/abnormal decision is based on the absence or presence of textural abnormalities in the image (see Section 2).

Lung segmentations are used to determine the initial symmetry axis and limit the computation of local symmetry. They were obtained using a previously developed algorithm, which is a combination of pixel classification and shape modeling.²⁸

The discriminatory power of S to distinguish normal and abnormal images was evaluated using the area A_z under the receiver operating characteristic (ROC). In this experiment, the influence of the algorithm parameters, namely patch size m , position weight w , and subsample factor κm was studied. The effect of varying parameter values was first determined per individual parameter; as starting values for each experiment we used $w = 10.0$, $m = 9$ pixels, $\kappa = 1$. These values were found to work well in a patch-based categorization and retrieval method involving chest radiographs by Avni *et al.*²² After this first approximation, the optimal combination of w and m was determined. The set of angles A used to determine α_{opt} was $\{-10.0, -9.5, \dots, 9.5, 10.0\}$ degrees. The set of horizontal displacements Δ used to estimate δ_{opt} was $\{-10, -8, \dots, 8, 10\}$ pixels. Additionally, an experiment was performed with $m = 1$, $w = \infty$, optimal rotation, and optimal x -location. These values correspond to the situation where symmetry is computed as the (normalized) difference between a pixel and its corresponding position mirrored pixel.

2.F.2. Local symmetry in combination with texture analysis

In this experiment, the effect of adding local symmetry to a set of texture features when classifying patches and images for the presence of textural abnormalities was evaluated.

The dataset is the same as used for the first experiment. For training of the patch classifier, labeled examples of patches are required. Manual outlines of abnormalities were created in the full set by an expert reader (LH) under supervision of an experienced chest radiologist. Patches whose center is inside the outline were assigned the label abnormal. Normal patches were only sampled from normal images. From the original images, 145,315 patches (116,252 normal and 29,063 abnormal) and from the optimally rotated images, 144,905 (115,924 normal and 28,981 abnormal) patches were extracted, both with a normal to abnormal ratio of exactly 4:1. Training and testing of the 348 cases was performed in twofold cross-validation, by dividing the data in two equally

sized sets of 174 images and alternating the role of training and test set.

Optimal parameter values (m, w) for local symmetry maps were computed in a similar setup as in the previous experiment, but only using the images in the training set. The image rotation and symmetry axis x -location were optimized and $k = 1$. The detection of textural abnormalities is based on texture analysis of circular image patches (radius = 32 pixels) sampled every 4 pixels. Texture features were computed by extracting statistics of Gaussian derivative filtered images of order 0 through 2 ($L, L_x, L_y, L_{xx}, L_{xy}, L_{yy}$), at scales 1, 2, 4, and 8 pixels. The first four moments (mean, standard deviation, skew, and kurtosis) of the intensity distribution of each Gaussian derivative, filtered image and the original image were computed from pixels inside the corresponding circular patch. This method has recently successfully been used to detect textural abnormalities related to TB in chest radiographs.^{31,32} Two general position features, namely the x - and y -position normalized to the image size, and four lung segmentation-derived position features, namely the x - and y -position normalized to the bounding box of the lung fields, the distance to the lung boundary, and the distance to the center of gravity of the lung fields, were added to the texture features. A total of 106 features per patch were extracted.³² Image patches were sampled inside the segmented lung fields and assigned an abnormality likelihood with a GentleBoost classifier³³ which used 100 regression stumps as weak classifiers. Image locations outside the mirror symmetric lung fields \mathbf{P} and \mathbf{P}' were assigned $s = 0$. Images were assigned an overall texture score by computing the 95th percentile of the cumulative distribution of patch likelihood scores.³² This texture score was used to determine image classification performance.

Two sets of features were compared: (a) the texture+position features totaling 106 features, and (b) the texture+position features and local symmetry totaling 107 features. These feature sets were compared in a patch classification and image classification experiment. In addition, the performance of local symmetry as a single feature was determined in the patch classification experiment. Classification performance was determined using A_z . Significant differences were determined with case-based bootstrapping³⁴ using 1000 bootstrap samples.

2.F.3. Local symmetry to detect nodules

Lung cancer is commonly detected on radiographs, but it is known that retrospectively visible lesions are missed by radiologists in 19–90% of cases,³⁵ so computerized support for lung nodule detection is an active area of research.³⁶ Some nodules are very difficult to detect, also by humans, and a careful comparison of the left and right lung fields is often required to discern them. As an illustration of the local symmetry algorithm, an experiment was performed to determine its potential in enhancing the contrast of nodules on CXRs. In this experiment, also the effect of preprocessing with local normalization was determined. Local

normalization improves contrast of nodules,²⁹ and also serves to reduce low-frequency intensity variations which are uninformative for the detection of small lesions.

For evaluation we used the publicly available JSRT database consisting of 93 normal cases and 154 abnormal cases.³⁷ Only abnormal cases were used in this experiment; each contained one nodule of which location and radius were available. Images were digitized 12-bit posterior–anterior CXR, scanned to a resolution of 2048 × 2048 pixels with an isotropic pixel resolution of 175 μm. Nodule sizes ranged from 5 to 60 mm (median = 15 mm) with five degrees of conspicuousness, from obvious, relatively obvious, subtle, very subtle, and extremely subtle. These ratings were assigned by a radiologist. Four cases, in which the nodule was located outside the lung fields, could not be used; thus, we had 150 cases available for analysis.

Images were resampled to a width of 512 pixels. Local symmetry maps were computed with the optimal parameter values determined in the previous experiment. The image rotation and symmetry axis *x*-location were optimized.

The visibility of a nodule was determined by its contrast with its neighboring background, using the Weber contrast

$$C = \frac{I_f - I_b}{I_b} \tag{11}$$

where I_f is the average intensity of the nodule region of interest (ROI) and I_b the average background intensity. The Weber contrast, also known as the Weber ratio, is a commonly used measure used in perception studies as it follows Weber’s law of psychophysics³⁸. The nodule ROI is defined by a circle centered at the nodule location and with a radius r obtained from the JSRT annotations. I_b is measured in a band enclosing the nodule ROI with a width of $0.5r$. The contrast was measured only inside the unobscured lung fields, which were automatically segmented using active shape models.²⁸

C was computed on four types of input images: the original image, the locally normalized (LN) image with $\sigma_{LN} = 16$ pixels,²⁹ and local symmetry maps computed from the original and locally normalized image. Note that C can be negative when the surroundings have higher values than the nodule ROI.

3. RESULTS

3.A. Global symmetry to discriminate between normal and abnormal images

Figure 4 shows the effect of optimal rotation on individual cases for the default parameter settings. By definition all scores are equal or lower after optimization. For normal cases (green), the scores are reduced more than for the abnormal cases (red), as indicated by their respective trendlines. This differential change between normal and abnormal cases leads to a large improvement in discriminative performance.

Figure 5 shows the effect of varying the free parameters on A_z for the 348 test images. Optimal rotation of the symmetry axis was important for the majority of parameter values

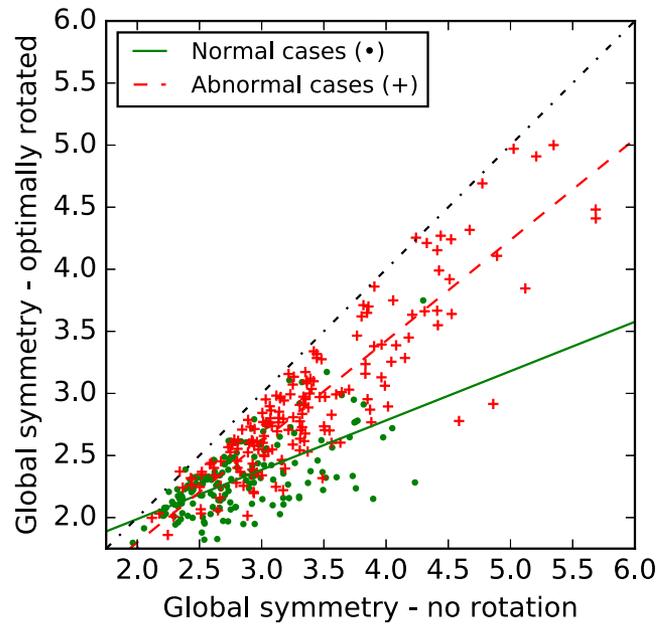


FIG. 4. Effect of optimal rotation on global symmetry S scores for 348 cases. Scores were computed with optimal parameter values. S for normal cases (dots) decreases relatively more than abnormal cases (pluses) as shown by the trendlines. As a result, projection on the horizontal axis (no rotation) results in high overlap of normal and abnormal cases, but projection on the vertical axis (optimal rotation) yields good separation. [Color figure can be viewed at wileyonlinelibrary.com]

and led to a large increase in performance. Additional optimization of the x -location did not lead to large further increases. To investigate the reason for this, we investigated the absolute changes in optimal x -location compared to the initial symmetry axis. The mean absolute change in x -location was 1.15 ± 1.38 pixels and the median absolute change was 0 pixels. These values indicate that the initial x -location of the symmetry axis was very close to the optimal location for most images.

Performances reported in the remainder of this section refer to the images with optimal rotation and x -location. A_z showed a slow increase with increasing m up to a value of 13 pixels and then a slow decrease. In the m range of 9–21 pixels, A_z values were highly stable. At $w = 5.62$, the optimal A_z of 0.835 was achieved. In the w range of 3.2–32, A_z was mostly stable with values >0.82 . Higher and lower values of w led to a reduction in performance. Especially for low values of w , in which case the influence of \mathbf{f}_{img} increases relatively to \mathbf{f}_{pos} , performance decreased substantially. The relation between κ and A_z was mostly stable for $\kappa \leq 16$, for $\kappa > 16$ performance slightly reduced until it breaks down at $\kappa = 256$. The experiment with $m = 1$ and $w = \infty$ resulted in $A_z = 0.732$, a value markedly lower than the optimal values in Fig. 5.

The value of w is related to the value m via Eq. (4); higher values of m result in a larger number of intensity features and therefore require higher values of w to maintain the same relative weighting between intensity and position features. To reflect this, all combinations of $m = (9, 11, 13, 15, 17)$ and $w = (1.77, 3.16, 5.62, 10.0, 17.7, 31.6)$ were tested with $\kappa = 16$ to determine the optimal combination. For these parameter

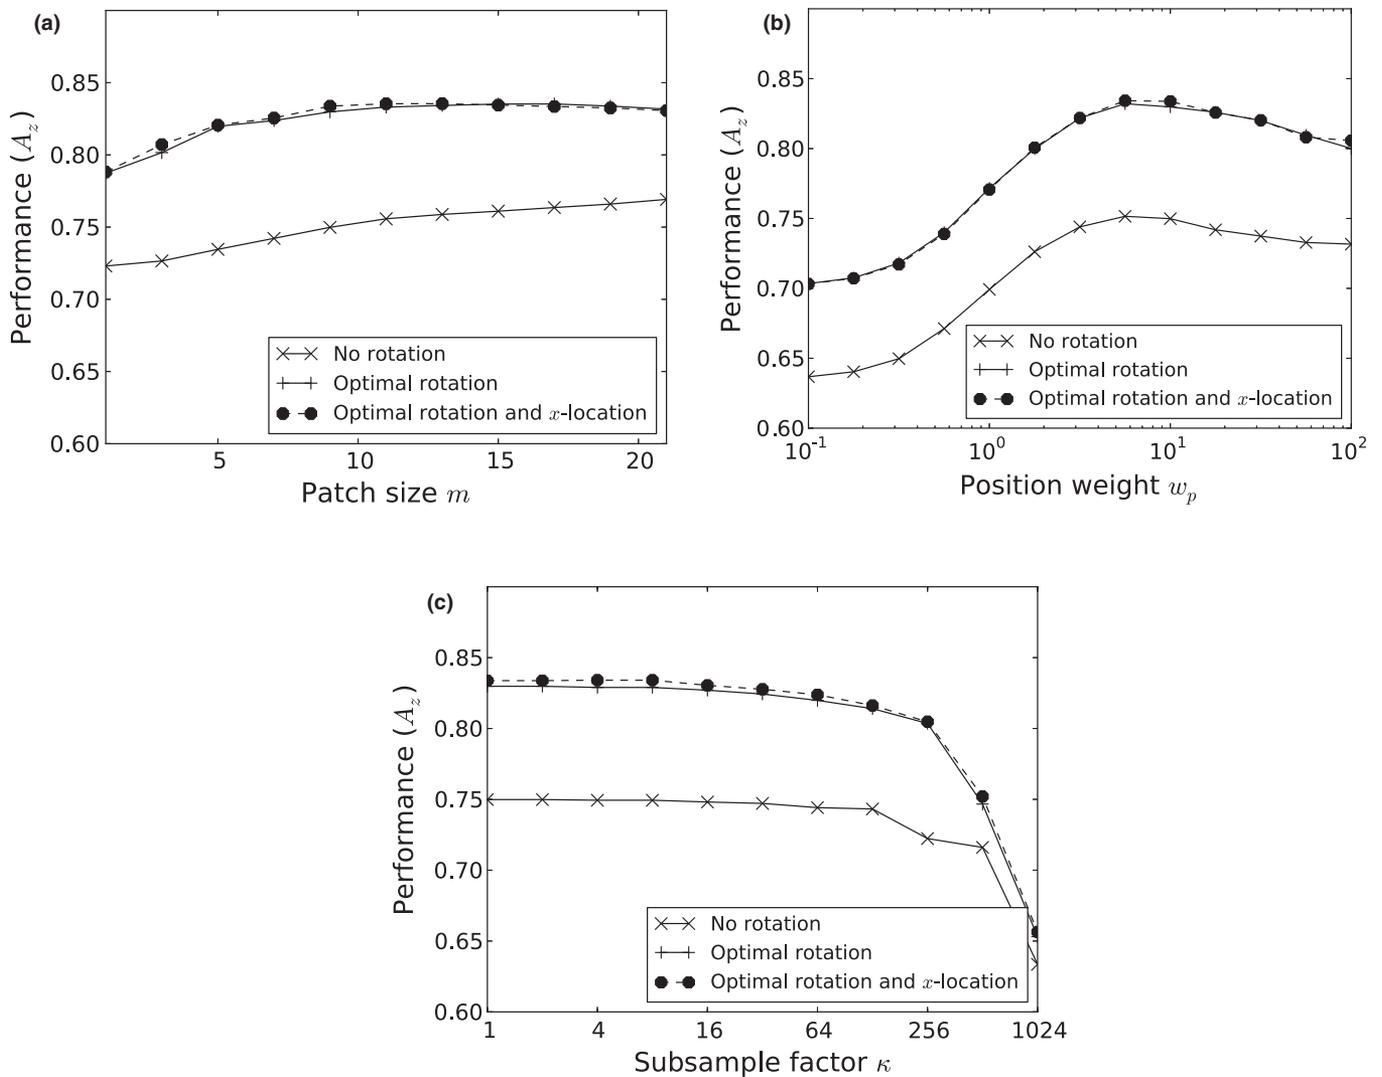


FIG. 5. Optimization of position weight w , subsample factor κ , and patch size m for global symmetry computation on a set of 348 images. Image classification performance for original images, optimally rotated images, and optimal x -location of the symmetry axis are shown as function of the three parameters: (a) Patch size, (b) Position weight (NB the x -axis is logarithmically scaled), (c) Subsample factor (NB the x -axis is logarithmically scaled).

values, the highest performance was found in their individual optimization. The optimal combination was $(m, w) = (15, 17.7)$ with $A_z = 0.838$; these parameter values were used in subsequent experiments.

Computation times are related to the amount of subsampling. Computation times for local and global symmetry combined (at a single core of a Core 2 Duo @ 3.0 GHz) decreased from 50 s ($\kappa = 1$) to 13 s ($\kappa = 4$) and 8 s ($\kappa = 16$) for an average case ($\pm 40,000$ positions), using optimization of image rotation and x_s , and with optimal (m, w) values. Subsampling with a factor $\kappa = 16$ gives minimal performance loss compared to $\kappa = 1$. An additional reduction of computation time could be achieved by not performing x_s optimization, as it has only a very small effect on performance. The search for optimal corresponding patches is also easy to parallelize in order to further improve the computation times.

Figure 6 shows an example of the effect of w on the computation of local symmetry. For values of w that are

low compared to the optimal range (0.1 and 1), i.e., a low influence of position in searching for corresponding points, matching can occur far from the position-symmetric point. In this case, even for abnormal patches a corresponding patch can be found somewhere in the contralateral side, resulting in overall high local symmetry values. In the case of values that are near the optimal range ($w = 10$), corresponding patches can usually be found for normal patches, but not for abnormal ones. This results in a discrimination between normal and abnormal patches. For high values ($w = 100$), the search space is effectively limited to the exact position-symmetric point. Although this results in low local symmetry values for some of the abnormal areas that remained inconspicuous with lower values of w , it also results in lower local symmetry values for normal areas. More importantly it is not possible to localize abnormalities anymore as the local symmetry map has become (largely) symmetric.

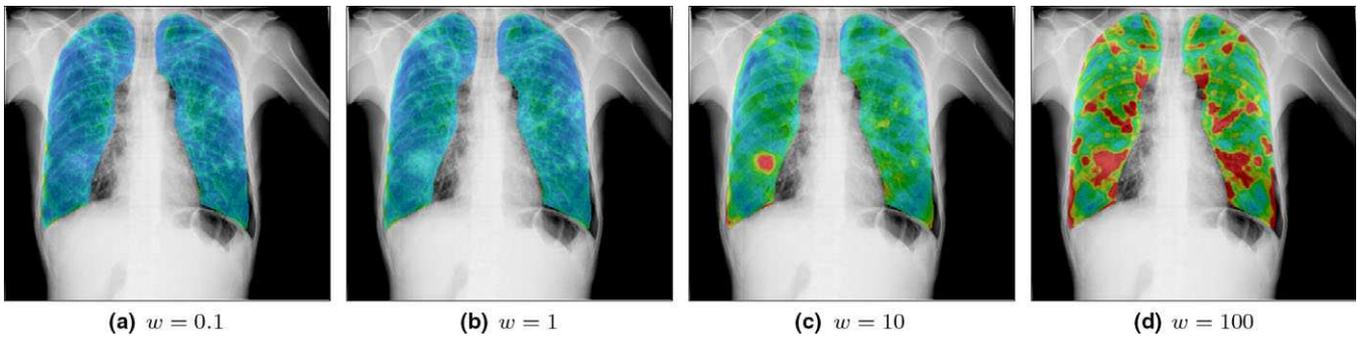


FIG. 6. The effect of the weighting factor w on the computation of local symmetry. Example for one image where local symmetry maps were computed with values of $w \in \{0.1, 1, 10, 100\}$. Overlay scaling is identical for all symmetry maps. [Color figure can be viewed at wileyonlinelibrary.com]

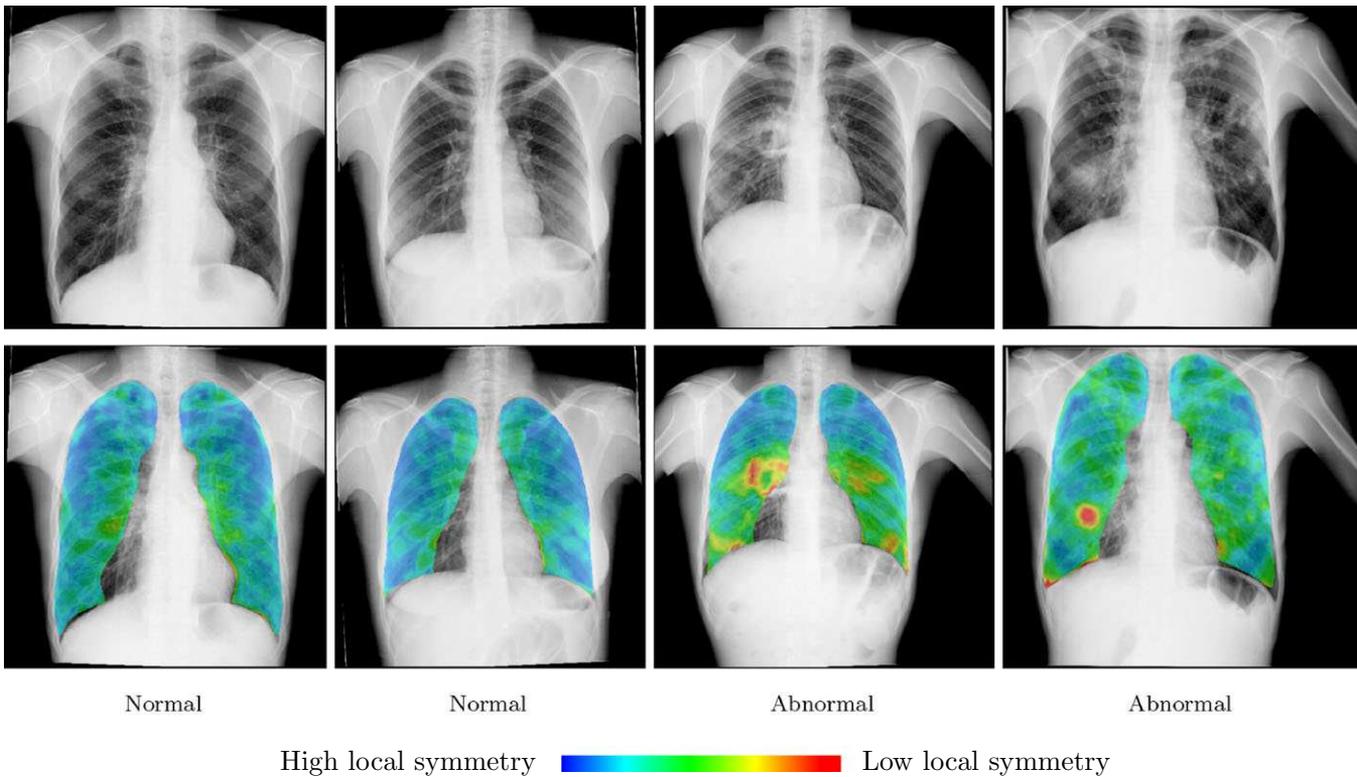


FIG. 7. Examples of local symmetry for two normal and two abnormal images. Overlay maps indicate local symmetry; overlay scale has arbitrary units, because s values have only a relative interpretation. The scaling is identical in all the examples. [Color figure can be viewed at wileyonlinelibrary.com]

3.B. Local symmetry in combination with texture analysis

Optimal parameter values for local symmetry computation were $(m,w) = (15, 17.7)$ for both the training sets, similar as in Section 3.A, but optimal A_z values were 0.843 and 0.828 for the two training sets, respectively. Figure 7 shows four examples of CXRs: the first two contain no textural abnormalities, the last two contain several abnormalities across the lung fields. Local symmetry maps are shown for all four cases. In the normal cases, values of s are on average low, with slightly higher values close to the hilar structures. In the

abnormal cases, abnormalities are highlighted in the local symmetry map.

Figure 8(a) shows the results of the patch classification experiment. Local symmetry as a single feature achieved $A_z = 0.726$. The texture+position features achieved $A_z = 0.878$. The addition of the local symmetry features to the texture+position features significantly increased performance to $A_z = 0.891$ ($P = 0.001$).

Figure 8(b) shows the results of the image classification experiment. The texture+position features achieved $A_z = 0.848$. The addition of the local symmetry features to the texture+position features significantly increased

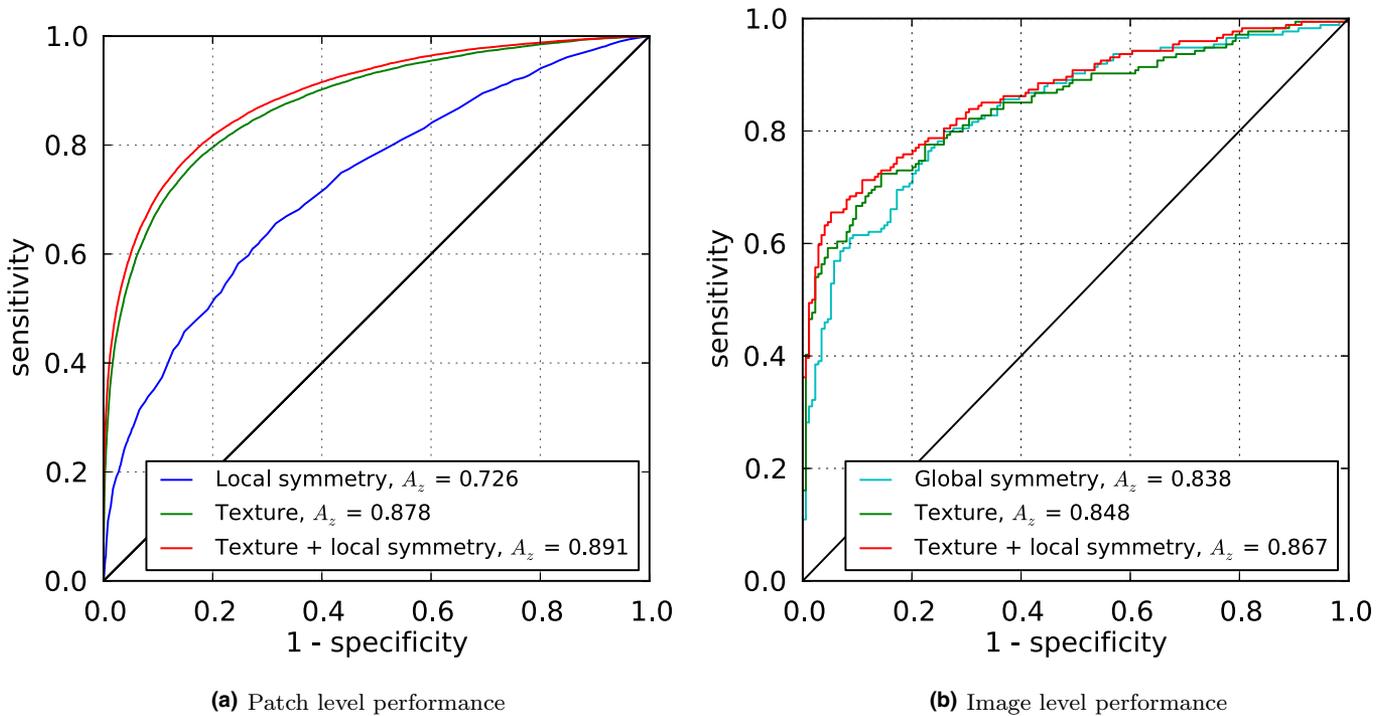


FIG. 8. Effect of adding local symmetry to a supervised system detecting textural abnormalities analyzed using ROC analysis. Experiments were done on 384 chest radiographs. *Texture* is the basic system without local symmetry, *Texture + local symmetry* includes *s* as a patch feature. (a) Patch (local) level performance. *Local symmetry* is the system with *s* as the only patch feature. (b) Image level performance. *Global symmetry* is added for reference and shows the performance of *S* as a single image feature. [Color figure can be viewed at wileyonlinelibrary.com]

performance to $A_z = 0.867$ ($P = 0.01$). For comparison, the results of using global symmetry to classify images is also shown in the figure. No significant difference was found in performance between global symmetry and the texture+position features alone ($P = 0.352$).

3.C. Local symmetry to detect nodules

Figure 9 shows three examples of original CXRs-containing nodules, locally normalized images, local symmetry maps, and the nodule ROI. In the local symmetry maps, an increase of the values relative to the surroundings is observed at the nodule locations. Note the near absence of rib and clavicle patterns, which are one of the most prominent structures in CXRs but do not show a pronounced response in the local symmetry map because they exhibit strong symmetry. At a number of locations in the image *s* is increased although no abnormalities are present. This can be observed near the pulmonary vessel tree, at some crossings of ribs and vessels and close to the lateral rib cage.

Figure 10(a) shows the average nodule *C* for original images and LN images and for local symmetry maps computed from original and LN images. The local normalization procedure on its own increased *C*, but local symmetry computation applied to the locally normalized image further enhanced the contrast. For both types of images, *C* increased significantly in the local symmetry map: from $1.3 \cdot 10^{-2}$ to $6.8 \cdot 10^{-2}$ ($P < 0.001$; paired Student's *t*-test) when using

original images and from $9.1 \cdot 10^{-2}$ to $14.1 \cdot 10^{-2}$ ($P < 0.001$; paired Student's *t*-test) when using LN images. For all of the subgroups of nodules, *C* consistently increased in the local symmetry maps, except for the extremely subtle nodules. These nodules are almost impossible to observe, also for human experts.³⁷

Figure 10(b) shows the changes in *C* per nodule computed on LN images and local symmetry maps computed from the LN image. For most cases, an improvement of *C* is observed. In LN images, the maximum contrast was $32.5 \cdot 10^{-2}$ and in local symmetry maps, $71.8 \cdot 10^{-2}$.

4. DISCUSSION

We have presented a method to compute local and global symmetry efficiently in 2D gray value images. In applications concerning detection of pathology in chest radiographs, global symmetry was found to be a strong indicator for the overall presence of abnormalities and local symmetry was an informative measure for localizing abnormalities. The method provides three contributions to the field of automatic medical image analysis: (a) a novel continuous symmetry measure was developed, (b) to our knowledge this is the first method that uses global symmetry to detect images containing abnormalities, (c) the method addresses the issue of inherent and pathological asymmetry by combining gray value and position information to quantify symmetry. The proposed method for computing symmetry can contribute to the improvement of automated abnormality detection systems for any organ

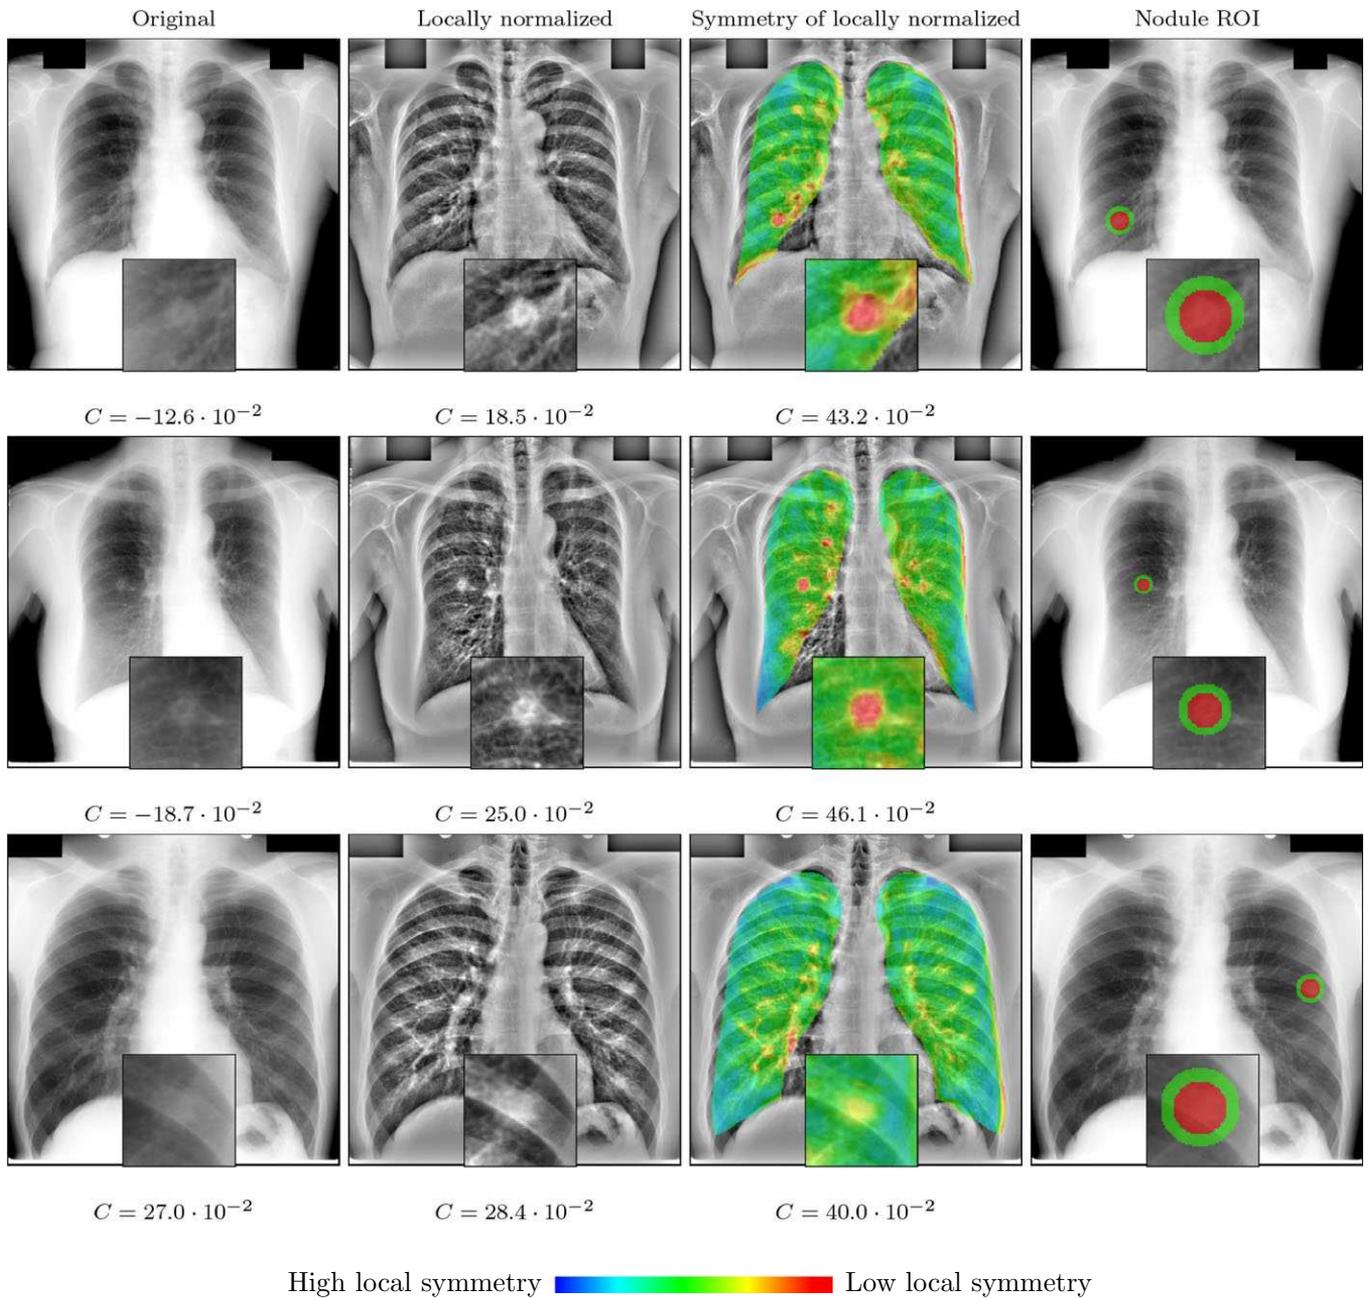


FIG. 9. Local symmetry maps for three CXRs-containing nodules. Shown are original images, locally normalized images, symmetry maps calculated from locally normalized images, and the nodule ROI (inside) and background region (outside) used for contrast computation. Insets show a detailed view of the nodule ROI. The nodule contrast C is indicated below the images. [Color figure can be viewed at wileyonlinelibrary.com]

and imaging modality where (partial) symmetry is present in the nonpathological situation.

A property of symmetry computation is that, without adding additional information, it leads to symmetrical structure in the symmetry map. Especially when the computation is limited to spatially symmetrical positions, such as in most previously published methods on symmetry computation for pathology detection,^{9,15,16} the symmetry map is by definition fully symmetric. In medical images, which often do not exhibit perfect symmetry even in normal examples, this leads to an artificial elevation of symmetry measures. It will also lead

to an ambiguity of the side where abnormalities, such as pathology, are located. We have addressed this important issue by allowing flexibility in matching positions, using a weighting factor which controls the influence of position and local density patterns. Smaller abnormalities can then be unambiguously localized, such as is visible in Fig. 9 where nodules are highlighted on the correct side in the local symmetry map or Fig. 6c for an example of a correctly localized textural abnormality. In contrast, using only the exact reflected point for symmetry computation is equivalent to employing very high position weights in the

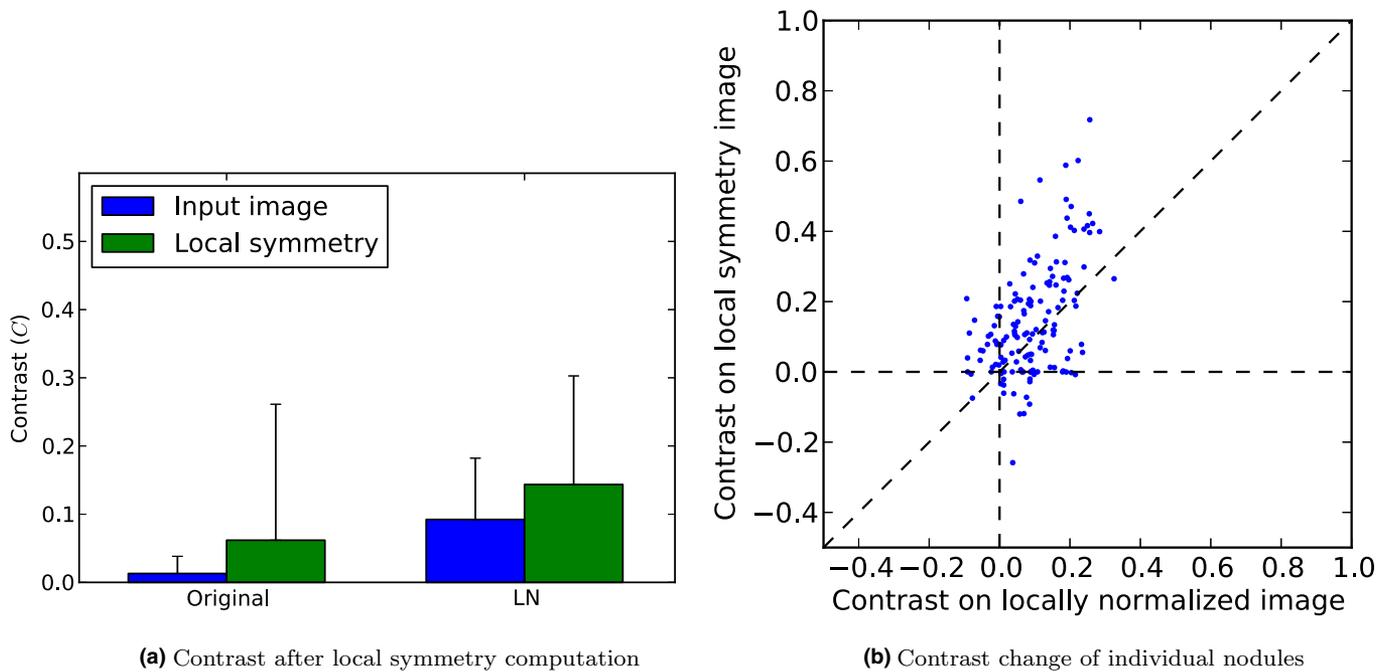


FIG. 10. Nodule contrast C measured in input images and local symmetry maps computed for 150 nodules. (a) Change of C measured in original and LN images. Error bars indicate standard deviation. (b) Change in C per nodule on locally normalized images and on local symmetry images computed on LN images. Each marker indicates a nodule. [Color figure can be viewed at wileyonlinelibrary.com]

proposed method (cf. Fig. 6c and 6d). Our results showed that these higher weights have lower discriminative performance than lower weights, indicating the importance of allowing positional flexibility in symmetry computation for pathology detection. Other applications or other modalities than chest radiographs likely require different optimal parameters, which need to be determined experimentally. As a general guideline, it is expected that higher values of w (favoring the position component) are possible when the type of image under consideration has a high degree of natural symmetry.

In the field of automatic CXR analysis, one of the few methods that uses left/right symmetry directly is the contralateral subtraction technique by Li *et al.*^{17,18} After determination of the symmetry axis based on rib profiles, the axis is rotated upright using a minimization procedure. Then global and elastic registration were used to align the lung structures. Although ribs were visually determined to have been removed in the majority of images, no evaluation was performed of the method's value in pathology detection. An important difference with our method is that we do not aim to solve the registration problem. Registration can be difficult, especially in pathological images, requires appropriate regularization, and leads to image artifacts.¹⁶ In fact, there is no perfect registration possible between contralateral lung fields. Instead our method derives diagnostic information from the absence of corresponding points between the right and the left lung and as such is less influenced by criteria required for proper image registration such as smoothness of the deformation field or inverse consistency between corresponding points.

One of the most outstanding results in this work is that global symmetry as a single unsupervised feature performed as well for detecting abnormal CXRs as a previously published supervised method, which analyzes CXR locally for abnormalities based on labeled training examples.^{31,32} Several reasons explain this result. The most important one is that in symmetry computation the CXR is used as its own reference for determining what is normal and abnormal; an important observation also mentioned by Sun *et al.*¹⁶ In this way, the problems of inherent (nonpathological) differences between CXRs from different individuals, due to anatomical (e.g., shape and appearance of the ribcage) and physiological (e.g., age), and also acquisition differences (e.g., scanner model), are addressed at the same time. This self-normalizing property holds for any type of medical image. A second reason is that global symmetry provides a robust holistic interpretation of the full radiograph without a need to determine what kind of patterns are present. In this sense, the method shows similarities to the first phase of the radiologist's reading process, in which a very short impression of the *Gestalt* of the image already provides a first clue to the presence of abnormalities.^{39,40} A disadvantage of using only global symmetry, which is a summary statistic of local symmetry, is that it can only detect images with relatively large abnormalities, therefore we also investigated the diagnostic properties of local symmetry.

We found that local symmetry was an informative feature, which improved detection of textural and small nodular abnormalities. When added to a set of texture features in a state-of-the-art supervised system, it improved the detection of textural abnormalities compared to texture features alone.

It might be surprising that adding local symmetry improved performance significantly in a combination with a large set (106) of other features. We hypothesize that a reason for this is the use of the nonlinear GentleBoost classifier. This type of classifier uses an implicit feature selection in each iteration of its training phase, where the feature is selected that minimizes the error for the current weighting of samples in the dataset using a (weak) regression stump classifier.³³ An illustration of the importance of local symmetry is that, in the patch classification experiment, the feature was selected first — indicating that it has the highest discriminatory performance of all the features — and in total in 6 of 100 regression stumps of GentleBoost. In the third experiment, we found that the local symmetry enhances the contrast of nodules and in some cases even strongly highlighted the correct location of the nodule. This property could be used in automated nodule detection, in addition with other features, to improve detection and classification.

In generic 2D images, a full search for the position and orientation of the symmetry axis is required as these parameters are unknown. Existing methods in literature have therefore included methods to find the axis as an integral part of the algorithm.^{9,15} In medical images, it is often possible to make a good initial approximation based on the prior knowledge of the image content. Instead of requiring a full search, we refine this initial approximation, which can be usually accurately found because the danger of being trapped in a local minimum is small. Such an approach shares similarities to the work of Liu *et al.*,¹⁴ who used minimization techniques to determine the midsagittal plane in pathological brain MRI. We found that an optimal rotation of the image led to large performance increases of global symmetry compared to the original image. The optimization of the x -location of the symmetry axis led to only a minor performance increase. This is explained by two observations. First, the algorithm is relatively insensitive to translations in the x -direction; minor position changes of corresponding patches will lead to overall slightly higher local symmetry values, but not to a loss of the discriminative properties of global symmetry. Second, the initial x -location of the symmetry axis determined from the lung mask is already close to the optimal location. To further improve computation speeds, optimization techniques like Powell's or Nelder-Mead's method could be employed to find the optimal rotation angle and x -coordinate of the symmetry axis, instead of employing two sequential one-dimensional searches as were used in this work.

In the case of chest radiographs, the initial symmetry axis is determined from the lung segmentation. The lung segmentation can be obtained through various fully automatic algorithms.^{28,41} The segmentation serves the other purpose of excluding inherently asymmetrical parts of the image; in the case of CXRs, the heart region and regularly the presence of gas in stomach or colon or an asymmetrical positioning of the arms. No information regarding symmetry is available in that region; for the discriminative properties of global symmetry this is an issue when abnormalities are only present in the excluded areas. In these excluded areas, we set local

symmetry values to a value of 0, in order to still allow local analysis by combining local symmetry with other local features. Alternatively, classification techniques dealing with missing values could be used.

There are a number of situations where the optimization of the symmetry axis can fail. The most prominent one is the presence of very large, unexpected, asymmetric structures, such as gross pathology. The optimization can then be trapped in a local minimum, for example, when the symmetry axis has been positioned in such a way that the pathology is aligned with a structure of similar (high) density on the other side. This failure to detect the correct upright position in abnormal images is not necessarily a problem, as global symmetry scores will remain high. A more difficult situation is the presence of density differences on the sides of the symmetry axis which are not caused by pathological processes. An example of such a situation is rotation of the rib cage around the caudocranial axis, which causes a slight intensity difference between left and right lung fields. This specific problem could be addressed by using contrast invariant descriptors for \mathbf{f}_{img} , but might lead to loss of performance in abnormal images where density differences play an important role. In future research, the symmetry algorithm could be expanded to recover caudocranial rotation angles by analyzing the (global) density differences between left and right lung fields. Another situation where the procedure might fail, and lead to elevated symmetry scores, is the presence of strong local (elastic) deformations between corresponding structures in left and right lungs, but such deformations themselves can suggest pathological processes.

The optimal type of descriptor for use in symmetry computation could be further investigated, see Mikolajczyk *et al.*⁴² for an overview of existing techniques. Of the point descriptors, SIFT descriptors are a popular point descriptor and they have been successfully employed in symmetry computation.⁹ In this work, we used raw patch values. This type of descriptor was shown to have similar performance as SIFT descriptors in a content-based image retrieval application for the detection of abnormalities in CXRs.²² One of the reasons why raw intensity values work well in chest radiographs is that rotation and scale invariance are not required. On the contrary, the orientation of certain anatomical structures, such as ribs, provide valuable information for symmetry computation. In other types of medical images, it might be beneficial to explore different types of point descriptors and distance measures. Regarding distance measures one can think of cross-correlation to provide contrast invariance, and mutual information for nonlinear intensity relations between similar patches.

The presented method computes symmetry in 2D gray-scale images, but it can be easily extended to N-D images, such as CT and MRI, and color images, such as retinal or microscopic pathology images. The only requirement is that an appropriate point descriptor and similarity measure is used. Radial symmetric structures can be addressed in the same framework by a slight modification of the algorithm. Because the method is designed to deal with the presence of normal asymmetry, structures do not have to exhibit perfect

symmetry to be suitable for analysis. In this paper, we have shown the relevance of symmetry calculation for automatic detection systems, and in future research on the symmetry calculation, it could be interesting to investigate the use of symmetry maps as a visual aid to human readers.

5. CONCLUSIONS

An efficient method to quantify local and global symmetry in medical images was presented. The method is designed to work under conditions of normal inherent asymmetry and pathology-induced asymmetry. In three experiments on chest radiographs, it was demonstrated that local and global symmetry are strong indicators for the presence of pathology.

ACKNOWLEDGMENTS

The authors thank Helen Ayles from the London School of Hygiene & Tropical Medicine and the Zambia AIDS-Related TB (ZAMBART) Project for providing the TB suspects dataset used in this study.

CONFLICTS OF INTEREST

Bram van Ginneken is cofounder and shareholder of Thirona. The other authors have nothing to disclose.

^{a)}Author to whom correspondence should be addressed. Electronic mail: bram.vanginneken@radboudumc.nl.

REFERENCES

- Saber E, Tekalp AM. Frontal-view face detection and facial feature extraction using color, shape and symmetry based cost functions. *Pattern Recognit Lett*. 1998;19:669–680.
- Javed O, Shah M. Tracking and object classification for automated surveillance. In: Heyden A, Sparr G, Nielsen M, Johansen P, eds. *Computer Vision – ECCV 2002. ECCV 2002. Lecture Notes in Computer Science*. Vol. 2353. Berlin, Heidelberg: Springer; 2002:343–357.
- Li WH, Kleeman L. Real time object tracking using reflectional symmetry and motion. In: *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, Beijing*. 2006:2798–2803.
- Hays J, Leordeanu M, Efros AA, Liu Y. Discovering texture regularity as a higher-order correspondence problem. In: Leonardis A, Bischof H, Pinz A, eds. *Computer Vision – ECCV 2006. ECCV 2006. Lecture Notes in Computer Science*. Vol. 3952. Berlin, Heidelberg: Springer; 2006:522–535.
- Liu Y, Hel-Or H, Kaplan CS, Gool LV. Computational symmetry in computer vision and computer graphics. *Found Trends Comp Graphics and Vision*. 2010;5:1–195.
- Giblin PJ, Kimia BB. On the intrinsic reconstruction of shape from its symmetries. *IEEE Trans Pattern Anal Mach Intell*. 2003;25:895–911.
- Marola G. On the detection of the axes of symmetry of symmetric and almost symmetric planar images. *IEEE Trans Pattern Anal Mach Intell*. 1989;11:104–108.
- Zabrodsky H, Peleg S, Avnir D. Symmetry as a continuous feature. *IEEE Trans Pattern Anal Mach Intell*. 1995;17:1154–1166.
- Loy G, Eklundh J-O. Detecting symmetry and symmetric constellations of features. In: Leonardis A, Bischof H, Pinz A, eds. *Computer Vision – ECCV 2006. ECCV 2006. Lecture Notes in Computer Science*. Vol. 3952. Berlin, Heidelberg: Springer; 2006:508–521.
- Scognamiglio R, Rhodes G, Morrone C, Burr D. A feature-based model of symmetry detection. *Proc R Soc Lond B Biol Sci*. 2003;270:1727–1733.
- Grammer K, Thornhill R, et al., Human (homo sapiens) facial attractiveness and sexual selection: the role of symmetry and averageness. *J Comp Psychol*. 1994;108:233–242.
- Graham JH, Raz S, Hel-Or H, Nevo E. Fluctuating asymmetry: methods, theory, and applications. *Symmetry*. 2010;2:466–540.
- Keinan S, Avnir D. Quantitative symmetry in structure-activity correlations: The near C2 symmetry of inhibitor/HIV protease complexes. *J Am Chem Soc*. 2000;122:4378–4384.
- Liu Y, Collins RT, Rothfus WE. Robust midsagittal plane extraction from normal and pathological 3-D neuroradiology images. *IEEE Trans Med Imaging*. 2001;20:175–192.
- Sun Y, Bhanu B. Reflection symmetry-integrated image segmentation. *IEEE Trans Pattern Anal Mach Intell*. 2012;34:1827–1841.
- Sun Y, Bhanu B, Bhanu S. Automatic symmetry-integrated brain injury detection in MRI sequences. In: *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, Miami, FL*. 2009:79–86.
- Li Q, Katsuragawa S, Doi K. Improved contralateral subtraction images by use of elastic matching technique. *Med Phys*. 2000a;27:1934–1942.
- Li Q, Katsuragawa S, Ishida T, et al. Contralateral subtraction: a novel technique for detection of asymmetric abnormalities on digital chest radiographs. *Med Phys*. 2000b;27:47–55.
- Yoshida H. Local contralateral subtraction based on bilateral symmetry of lung for reduction of false positives in computerized detection of pulmonary nodules. *IEEE Trans Biomed Eng*. 2004;51:778–789.
- Kowner R, Thornhill R. The imperfect organism: on the concept of fluctuating asymmetry and its significance in human, non-human animals, and plants. *Symmetry: Culture and Science*. 1999;10:227–244.
- Good CD, Johnsrude I, Ashburner J, Henson RN, Friston KJ, Frackowiak RS. Cerebral asymmetry and the effects of sex and handedness on brain structure: a voxel-based morphometric analysis of 465 normal adult human brains. *Neuroimage*. 2001;14:685–700.
- Avni U, Greenspan H, Konen E, Sharon M, Goldberger J. X-ray categorization and retrieval on the organ and pathology level, using patch-based visual words. *IEEE Trans Med Imaging*. 2011;30:733–746.
- Lowe DG. Object recognition from local scale-invariant features. In: *Proceedings of the Seventh IEEE International Conference on Computer Vision, Kerkyra*. Vol. 2. 1999:1150–1157.
- Sivic J, Zisserman A. Video google: a text retrieval approach to object matching in videos. In: *Proceedings Ninth IEEE International Conference on Computer Vision, Nice, France*. Vol. 2. 2003:1470–1477.
- Ojala T, Pietikainen M, Maenpaa T. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans Pattern Anal Mach Intell*. 2002;24:971–987.
- Liu T, Moore AW, Gray A, Yang K. An investigation of practical approximate nearest neighbor algorithms. *Adv Neur In*. 2004;17:825–832.
- Arya S, Mount DM, Netanyahu NS, Silverman R, Wu AY. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *J ACM*. 1998;45:891–923.
- van Ginneken B, Stegmann MB, Loog M. Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database. *Med Image Anal*. 2006;10:19–40.
- Schilham AMR, van Ginneken B, Loog M. A computer-aided diagnosis system for detection of lung nodules in chest radiographs with an evaluation on a public database. *Med Image Anal*. 2006;10:247–258.
- Hogeweg L, Sánchez CI, de Jong PA, Maduskar P, van Ginneken B. Clavicle segmentation in chest radiographs. *Med Image Anal*. 2012;16:1490–1502.
- Arzhaeva Y, Hogeweg L, de Jong PA, Viergever MA, van Ginneken B. Global and local multi-valued dissimilarity-based classification: application to computer-aided detection of tuberculosis. In: Yang GZ, Hawkes D, Rueckert D, Noble A, Taylor C, eds. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009. MICCAI 2009. Lecture Notes in Computer Science*. Vol. 5762. Berlin, Heidelberg: Springer; 2009:724–731.

32. Hogeweg L, Mol C, de Jong PA, Dawson R, Ayles H, van Ginneken B. Fusion of local and global detection systems to detect tuberculosis in chest radiographs. In: Jiang T, Navab N, Pluim JPW, Viergever MA, eds. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2010. MICCAI 2010. Lecture Notes in Computer Science*. Vol. 6363. Berlin, Heidelberg: Springer; 2010:650–657.
33. Friedman J, Hastie T, Tibshirani R. Special invited paper. additive logistic regression: a statistical view of boosting. *Ann Stat*. 2000;28:337–374.
34. Rutter CM. Bootstrap estimation of diagnostic accuracy with patient-clustered data. *Acad Radiol*. 2000;7:413–419.
35. Quekel LG, Kessels AG, Goei R, Engelshoven JMV. Miss rate of lung cancer on the chest radiograph in clinical practice. *Chest*. 1999;115:720–724.
36. de Boo DW, Prokop M, Uffmann M, van Ginneken B, Schaefer-Prokop CM. Computer-aided detection (CAD) of lung nodules and small tumours on chest radiographs. *Eur J Radiol*. 2009;72:218–225.
37. Shiraishi J, Katsuragawa S, Ikezoe J, et al. Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *Am J Roentgenol*. 2000;174:71–74.
38. Gescheider G. *Psychophysics: The Fundamentals*. Mahwah, NJ: L. Erlbaum Associates; 1997.
39. Kundel HL, Nodine CF. Interpreting chest radiographs without visual search. *Radiol*. 1975;116:527–532.
40. Koontz NA, Gunderman RB. Gestalt theory: implications for radiology education. *AJR Am J Roentgenol*. 2008;190:1156–1160.
41. Seghers D, Loeckx D, Maes F, Vandermeulen D, Suetens P. Minimal shape and intensity cost path segmentation. *IEEE Trans Med Imaging*. 2007;26:1115–1129.
42. Mikolajczyk K, Schmid C. A performance evaluation of local descriptors. *IEEE Trans Pattern Anal Mach Intell*. 2005;27:1615–1630.