

Reduction of Lymph Tissue False Positives in Pulmonary Embolism Detection

Bernard Ghanem^{a,b}, Jianming Liang^a, Jinbo Bi^a, Marcos Salganicoff^a, and Arun Krishnan^a

^aSiemens Medical Solutions USA, Inc, 51 Valley Stream Parkway, Malvern, PA

^bUniversity of Illinois at Urbana-Champaign, 405 North Mathews Avenue, Urbana, IL

ABSTRACT

Pulmonary embolism (PE) is a serious medical condition, characterized by the partial/complete blockage of an artery within the lungs. We have previously developed a fast yet effective approach for computer aided detection of PE in computed topographic pulmonary angiography (CTPA),¹ which is capable of detecting both acute and chronic PEs, achieving a benchmark performance of 78% sensitivity at 4 false positives (FPs) per volume. By reviewing the FPs generated by this system, we found the most dominant type of FP, roughly one third of all FPs, to be lymph/connective tissue. In this paper, we propose a novel approach that specifically aims at reducing this FP type. Our idea is to explicitly exploit the anatomical context configuration of PE and lymph tissue in the lungs: a lymph FP connects to the airway and is located outside the artery, while a true PE should not connect to the airway and must be inside the artery. To realize this idea, given a detected candidate (i.e. a cluster of suspicious voxels), we compute a set of contextual features, including its distance to the airway based on local distance transform and its relative position to the artery based on fast tensor voting and Hessian “vesselness” scores. Our tests on unseen cases show that these features can reduce the lymph FPs by 59%, while improving the overall sensitivity by 3.4%.

Keywords: detection, X-ray CT, feature extraction, classification and classifier design, segmentation, tensor voting, multiple instance learning

1. INTRODUCTION

Pulmonary embolism (PE) is the third most common cause of death in the US with at least 650,000 cases occurring annually. PE is a sudden blockage in a pulmonary artery, and is caused by an *embolus* that is usually formed in the legs and travels in the bloodstream through the heart before reaching the lungs. PE is a very serious condition that can cause sudden death in about one-third of the cases. Most of those who die do so within 30 to 60 minutes after symptoms start. Anti-clotting medications are highly effective in treating PEs, but sometimes can lead to subsequent hemorrhage and bleeding. Therefore, they should be only given to those who really need. A major clinical challenge, particularly in an ER (Emergency Room) scenario, is to quickly and correctly diagnose patients with PE and then send them on to treatment – a prompt and accurate diagnosis is the key to survival.

However, PE is among the most difficult conditions to diagnose because its primary symptoms are vague, non-specific, and may have a variety of other causes, making it hard to separate out the critically ill patients suffering from PE. The diagnosis of PE is missed more than 400,000 times in the US each year, and approximately 100,000 patients die who would have survived with the proper diagnosis and treatment.

Computed tomographic (CT) pulmonary angiography (CTPA) has become first-line diagnosis technique for PE. Significant PE’s are detectable given the high spatial resolution of modern CT scanners. A CT image is a large 3D volumetric image, which consists of hundreds of images, each representing one slice of the lung. Clinically, manual reading of these slices is laborious, time consuming and complicated by various PE look-alikes (false positives or FPs) including respiratory motion artifacts, flow-related artifacts, streak artifacts, partial volume artifacts, stair step artifacts, lymph nodes, and vascular bifurcation, among many others. The accuracy and efficiency of interpreting such a large image data set is also limited by human factors, such as attention span

Further author information: (Send correspondence to bghanem2@uiuc.edu)

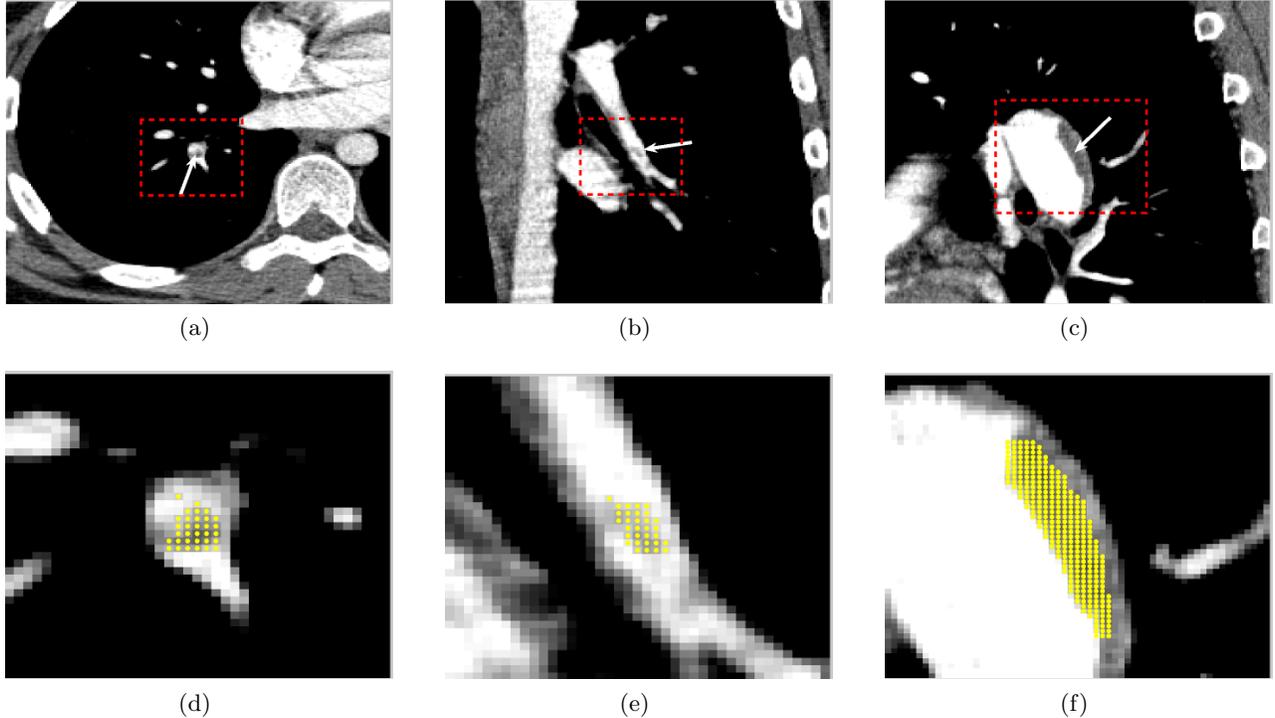


Figure 1. The emboli appears as dark regions residing in bright vessel lumen. The toboggan-based approach^{1,6} is able to detect both acute (a, b) and chronic (c) pulmonary emboli, offering simultaneous detection and segmentation (d, e, f). The clot in (b) was actually missed by our radiologists, but correctly detected by our system, and confirmed by the radiologists.

and eye fatigue. Consequently, it is highly desirable to have a computer aided detection (CAD) system to assist radiologists in detecting and characterizing emboli in an accurate, efficient and reproducible way. Such a CAD system has to achieve an extremely high detection sensitivity with as few as false positives to acquire clinical acceptance. It also needs to satisfy stringent real-time requirement due to the emergency nature of PE cases.

A number of computer aided diagnosis methods have been developed.²⁻⁵ These existing methods are all based on sophisticated vessel segmentation, namely, first segmenting the pulmonary vessel structure and then searching for PEs within the segmented vessels, because PEs only exist in pulmonary arteries. However, vessel segmentation is computationally time-consuming and has been problematic in small vasculature where subsegmental PEs often occur.² Furthermore, the normal regions of pulmonary vessels are enhanced with contrast material. There is no need to search for PE in the enhanced normal regions. Therefore, even if the pulmonary vascular structure is correctly segmented, a large part of it would be excluded anyway. In this context, we refer to our previous work,^{1,6} which uses a fast yet effective toboggan method embedded in a multiple instance learning framework to localize PEs in CT images without performing vessel segmentation. This original system is capable of detecting both acute and chronic PEs, achieving a benchmark performance of 78% sensitivity at 4 false positives (FPs) per volume. Figure 1 shows an example of the detection results of this system.

Despite their merits, the detection performance of the previous methods suffer from certain FP types especially those that share similar intensity distributions with true PEs (e.g. lymph/connective tissue). In fact, by reviewing the FPs generated by our original system,¹ we found the most dominant type of FP, roughly one third of all FPs, to be of the lymph type. In this paper, we propose a novel approach that specifically aims at reducing this FP type. Our idea is to explicitly exploit the anatomical context configuration of PE and lymph tissue in the lungs: a lymph FP connects to the airway and is located outside the artery, while a true PE should not connect to the airway and must be inside the artery. To realize this idea, given a detected candidate (i.e. a cluster of suspicious voxels), we compute a set of contextual features, including its distance to the airway

based on local distance transform and its relative position to the artery based on fast tensor voting and Hessian “vesselness” scores. In summary, the contributions of this work are twofold. **(i)** Novel features are developed to specifically discriminate lymph candidates from PEs. These features are based on anatomically grounded contextual cues that characterize lymph. **(ii)** We use a tensor voting framework that allows the accumulation of local information surrounding candidate clusters into global features. The main advantages of this frameworks includes its robustness to noise, invariance to translation, rotation, and reasonable scale change, as well as its independence of an underlying probabilistic model.

This paper is organized as follows. **Section 2** gives a detailed description of the proposed approach, where the lymph-specific features are illustrated in **Section 2.1-2.2** and the proposed classification scheme is portrayed in **Section 2.3**. We validate the experimental performance of this system in **Section 3**.

2. OVERVIEW OF PROPOSED APPROACH

Before proceeding with the detailed description of our proposed approach, we give a formal definition of the type of lymph tissue, whose FP rate we aim to reduce.

DEFINITION 1. *Lymph/connective tissue, in this paper, includes all tissue or tissue-like matter, which is in direct contact with both an airway and an artery, with intensity distribution very similar to that of true PEs.*

Guided by this broad definition of lymph FP, we propose a number of features that encode a candidate’s contextual information, in terms of two major anatomical landmarks, namely airways and arteries. Figure 2 shows the overall layout of our proposed system. In brief, the airway tract in a given CTPA volume is first segmented and then contextual features are extracted for each candidate by estimating its distance from the airway and its relative position to the nearest artery. In the following sections, we will describe each of the novel stages in detail.

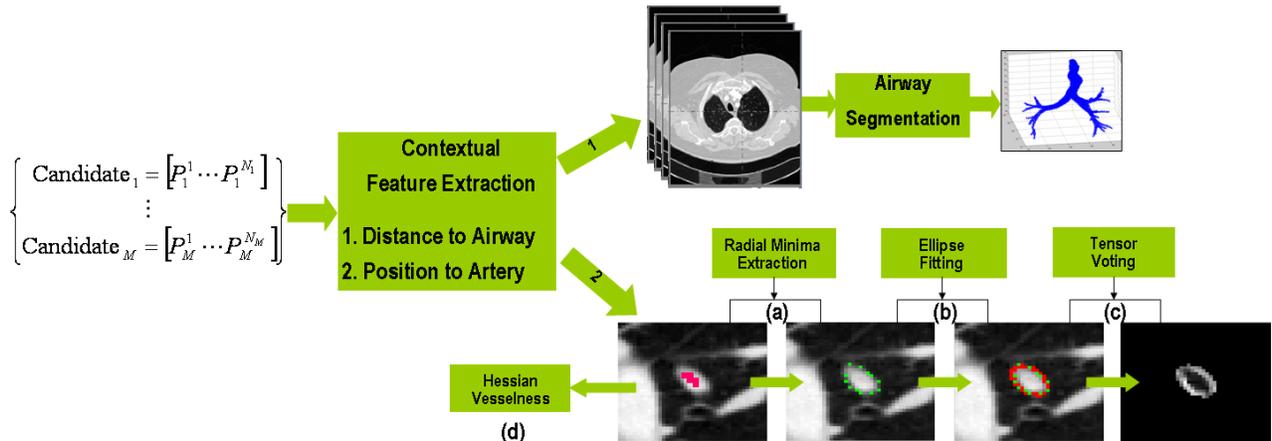


Figure 2. Overall layout of our proposed system.

2.1 Distance to Airway

In order to calculate the distance between the candidate and the airway, we first segment the airway tract, based on a local region growing algorithm. An initial seed point is chosen inside the airway tract (e.g. the carina position), from which segmentation commences. This is a recursive process, where a seed is grown to its direct neighbors according to decisions cast by termination templates, which guarantee propagation through the airway tract, while avoiding “leakage” at sites where the airway walls get thinner or disappear. Some of these templates use intensity based thresholding to cast their decisions. Here, we exploit the characteristic intensity distribution of airway voxels within CTPA volumes of the lungs (i.e. very dark regions). The other templates ensure that the grown region remains within the airway, by casting decisions on whether a local airway boundary has been reached. We spatially adapt the thresholds used by these templates and the size of the neighborhood around the seed point. In fact, fewer voxels are grown and a larger neighborhood is used at potential leakage points, as

compared to the interior of the airway. In this way, the growing method is guaranteed to converge to an accurate non-conservative segmentation. In summary, this simple algorithm combines the merits of local (bottom-up) airway segmentation techniques⁷ with those of global (top-down) techniques⁸, while significantly avoiding the “leakage” proneness of local methods and the computational expense and conservatism of global methods.

Figure 3 (a)-(c) show the results of applying this segmentation to three CTPA volumes of the lungs. After obtaining this segmentation, the distance between each candidate to the segmented airway is computed by performing a local 3D distance transform. In that way, the candidate-to-airway distance is computed as the minimum among all distances between the candidate’s voxels and the airway.

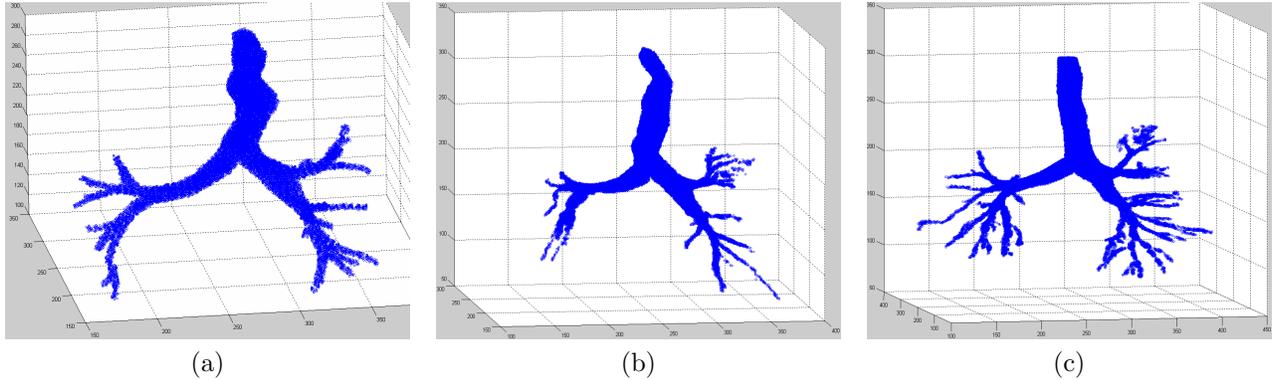


Figure 3. Three examples of segmentation produced by our region growing method, each sub-sampled five times in each dimension for visualization purposes.

2.2 Relative Position to Artery

In order to determine the relative position of a candidate to the artery, one may segment the artery and then detect whether the candidate is inside or outside the artery. However, this segmentation has proven to be time consuming and inaccurate especially in the presence of PEs.^{2,3,9} Hence, we resort to a local description of the context of a candidate relative to an artery. This is done using two sets of features: tensor voting features extracted using ellipse fitting in addition to “vesselness” features based on the Frangi, Sato, and Lorenz methods, as described below.

2.2.1 Tensor Voting Features

For each candidate, we first extract various cutting planes about its location, in each of which a set of features is extracted according to the general layout depicted in Figure 4. For a PE candidate, there exists a cutting plane, in which the candidate lies within an ellipse, representing the contours of an artery. For lymph candidates, no such plane exists. In each cutting plane, we determine intensity minima along each ray cast from the candidate (a), fit these minima to an elliptical model using eigen-fitting¹⁰ (b), and estimate the likelihood of whether this fitted ellipse represents an artery cross section, based on tensor voting (c).

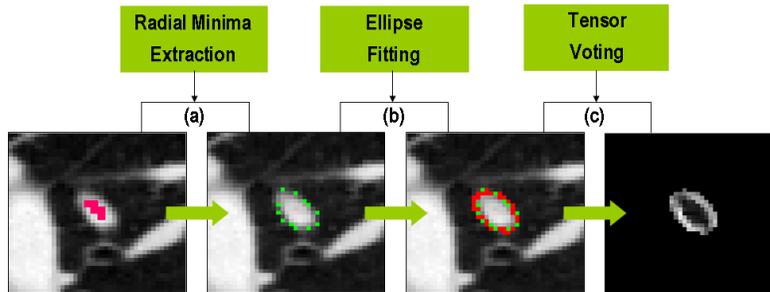


Figure 4. The general layout of extracting contextual features based on elliptical fitting and tensor voting.

Tensor voting is used here to evaluate the probability of the candidate lying within an artery. This is done non-parametrically without any predefined probability model. This method allows for the accumulation of local information (i.e. gradient direction and magnitude) at the local sites of the fitted ellipse to render global knowledge of its context, while remaining robust to noise and outliers. Here, we give a brief description of how tensor votes are computed and how the voting process is executed.^{11,12} For each cutting plane, we generate two sets of pixels: voter points, which are determined to be the points with high gradient energy in the plane, and receiver points, which are set to be the pixels constituting the interior and boundary of the fitted ellipse. Each voter point casts a separate second order tensor vote on each receiver point. This tensor vote is basically a rank-1 positive semi-definite matrix dependent on both the orientation of the voter with respect to the receiver (θ), the distance between them (l), and the stick tensor at the voter whose magnitude is denoted as $|S_V|$. In our setup, a voter’s stick tensor is set to its intensity gradient. In Figure 5(a), we show the geometric layout of how the second order vote is cast. In Figures 5(b)-(c), we plot the decay function of a voter at the origin with $|S_V| = 1$ and each second order vote cast by this voter to every receiver, respectively. The tensor vote between voter V and receiver R is computed in Equation (1). The only free parameter to be set here is the scale of the decay function defined by σ , which is set based on the size of the fitted ellipse.

$$s = \frac{l\theta}{\sin(\theta)}, \quad \kappa = \frac{2\sin(\theta)}{l} \Rightarrow DF(s, \kappa, \sigma) = |S_V|e^{-\frac{s^2 + c\kappa^2}{\sigma^2}}$$

$$TV(V \rightarrow R) = DF(s, \kappa, \sigma) \begin{bmatrix} \sin(2\theta)^2 & -\sin(2\theta)\cos(2\theta) \\ -\sin(2\theta)\cos(2\theta) & \sin(2\theta)^2 \end{bmatrix} \quad (1)$$

$$TV_{CUM}^R = \sum_{i=1}^{N_V} TV(V_i \rightarrow R) \quad (2)$$

The cumulative vote (TV_{CUM}^R) received at a receiver point is the sum of all the second order tensors cast by each of the N_V voter points to this receiver point. The tensor voting score (δ_R) for this point is the difference of the positive eigenvalues of TV_{CUM}^R . The higher δ_R is, the higher the probability that the receiver point lies on the boundary of a smooth curve. Also, the lower δ_R is, the higher the probability that the receiver point is within a smoothly bounded region. Figure 6 shows how the tensor voting score varies between true PE candidates and lymph FPs respectively. In case of a true PE candidate, there will be a cutting plane, with significantly low fitting cost and low voting score in the interior of the fitted ellipse but high voting score at the boundary of the ellipse. On the other hand, for candidates detected within lymph tissue, the fitting cost and the interior voting score will be much higher, while the boundary voting score will be much lower. Here, we note that to the authors’ best knowledge, the use of a tensor voting scheme to detect contextually based objects, such as pulmonary embolisms, is novel to medical imaging applications.

2.2.2 Hessian Based Vesselness Features

Our use of this type of feature is based on the basic idea that if a voxel is inside a vessel, then the Hessian matrix formed from second derivatives of intensity in the neighborhood of this voxel will have two negative eigenvalues and a third that is close to zero. The eigenvectors corresponding to the first two eigenvalues lie in a plane orthogonal to the central axis of the vessel, while the eigenvector corresponding to the third eigenvalue is in the direction of this central axis. Such discriminative power is manipulated to eliminate FPs lying outside blood vessels, namely lymph tissue. However, this only works if the vessel is empty of PEs. So, we precede this procedure by filling the detected region with high intensity values, to simulate the inside of a PE-ridden artery. To compute these Hessian features, we require the following steps:

- Fill the detected region with high intensity values
- Compute eigenvalues for each voxel in the candidate by:
 - Constructing the Hessian matrix

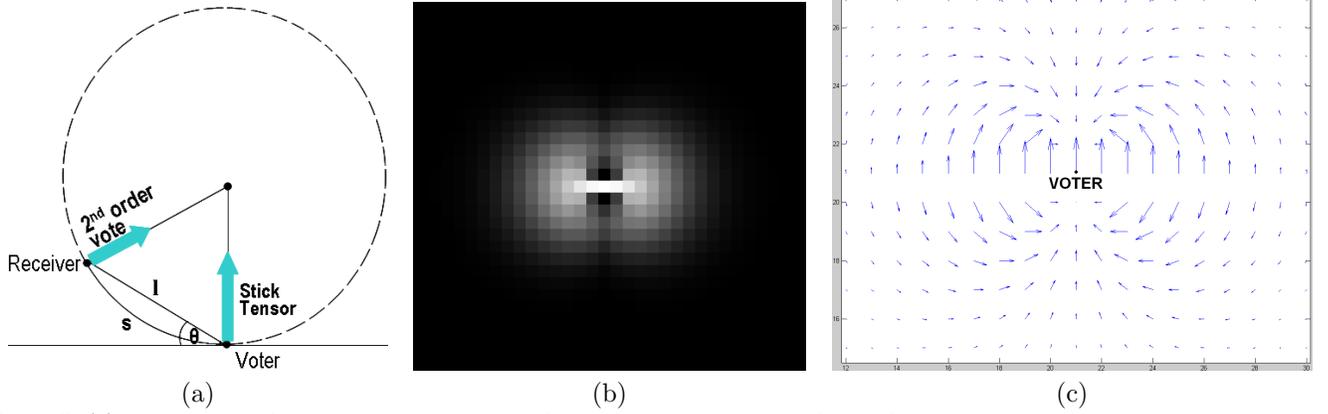


Figure 5. (a) shows the voting scheme of a voter casting a single vote at a receiver point, where the voting parameters, s , l , and θ are defined. (b) plots the decay function of a voter found at the origin with $|S_V| = 1$. (c) shows the second order tensor votes cast by a voter at the origin to every other point.

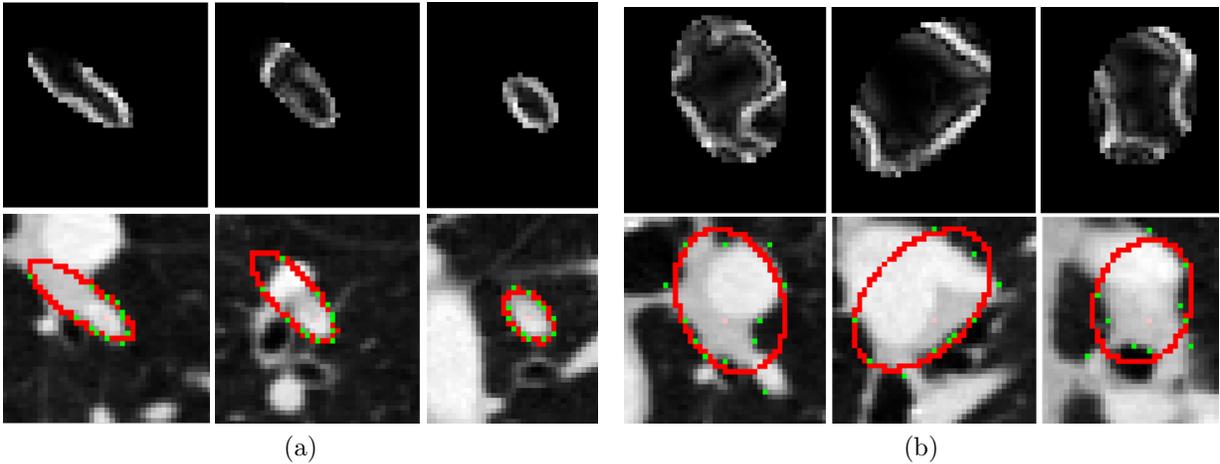


Figure 6. Tensor voting features extracted at the ellipse boundary and interior, for a true PE (a) and a lymph FP (b), each along three cutting planes.

- Computing the three eigenvalues and their ratios where $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$
- Computing three vesselness properties at each voxel (\vec{x}), based on the Frangi (3), Sato (4), and Lorenz (5) methods, with $\alpha = \beta = \gamma = \sigma = \xi = \tau = \rho = 1$:

$$\mathcal{F}(\vec{x}) = \left[1 - e^{-\frac{1}{2} \left(\frac{\lambda_2}{\alpha \lambda_3} \right)^2} \right] e^{-\frac{1}{2|\lambda_2 \lambda_3|} \left(\frac{\lambda_1}{\beta} \right)^2} \left[1 - e^{-\frac{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}{2\gamma^2}} \right] \quad (3)$$

$$\mathcal{S}(\vec{x}) = \begin{cases} \sigma^2 |\lambda_3| \left(\frac{\lambda_2}{\lambda_3} \right)^\xi \left(1 + \frac{\lambda_1}{|\lambda_2|} \right)^\tau, & \lambda_3 \leq \lambda_2 \leq \lambda_1 \leq 0 \\ \sigma^2 |\lambda_3| \left(\frac{\lambda_2}{\lambda_3} \right)^\xi \left(1 - \rho \frac{\lambda_1}{|\lambda_2|} \right)^\tau, & \lambda_3 \leq \lambda_2 \leq 0 \leq \lambda_1 \leq \frac{|\lambda_2|}{\rho} \end{cases} \quad (4)$$

$$\mathcal{L}(\vec{x}) = \sigma^\eta \left| \frac{\lambda_2 + \lambda_3}{2} \right| \quad (5)$$

- Compute statistics (i.e. maximum, minimum, mean, median and standard deviation) based on the computed eigenvalues and ratios for all the voxels in the candidate.

In the case of true PE, these vesselness features will be much higher in magnitude than those extracted in the case of lymph FP.

2.3 PE Classification

We design a multiple-instance classification algorithm based on Fisher’s linear discriminant (FLD) analysis.¹³ This multiple-instance learning approach makes use of the fact that some candidates can belong to the same PE, so that a PE is missed if and only if all the candidates belonging to it are detected as non-PE. Let the i^{th} PE contain n_i candidates, represented as a set of feature vectors $\{\vec{x}_{ij}\}_{j=1}^{n_i}$. Let S_i be the index set of all candidates that belong to the i^{th} PE. For each PE, we form a convex hull using these vectors in the feature space. Any point in the convex hull can be represented as a convex combination of \vec{x}_{ij} (i.e. $\sum_{j \in S_i} \lambda_{ij} \vec{x}_{ij}$ where $\lambda_{ij} \geq 0, \sum_{j \in S_i} \lambda_{ij} = 1$). The goal of our classification algorithm is to determine a decision boundary that separates, with high accuracy, any possible part of each of the convex hulls on one side and as many as possible negative detections on the other side. It implies that we do not require the entire convex hull to be correctly classified, but only any possible part of it. In other words, our algorithm solves the optimization problem in Equation (6), based on FLD analysis. C^+ and C^- are respectively the sets of positive (PE) and negative (non-PE) samples. y_i denotes the label, ξ_i is a residual error of the model fitting, $\|\vec{w}\|_2^2$ is the so-called regularization term that controls the classifier complexity, and γ represents the trade-off between the residual error and the complexity regularization. The classifier obtained by solving Equation (6) can dramatically reduce false detections as compared to standard algorithms, including FLD.

$$\begin{aligned} & \min \sum_{i=1}^m \xi_i^2 + \gamma \|\vec{w}\|_2^2 & (6) \\ \text{subject to } & \begin{cases} \vec{w}^T \left(\sum_{j \in S_i} \lambda_{ij} \vec{x}_{ij} \right) + b = y_i + \xi_i, \lambda_{ij} \geq 0, \sum_{j \in S_i} \lambda_{ij} = 1, \forall i \in C^+ \\ \vec{w}^T \vec{x}_i + b = y_i + \xi_i, \forall i \in C^-, \\ \sum_{i \in C^+} \xi_i = 0, \sum_{i \in C^-} \xi_i = 0 \end{cases} \end{aligned}$$

Since only a limited size of patient data is available, we resort to aggregation of multiple classifiers to reduce the variance of the learned classifier over various sample patient sets, thus enhancing accuracy. We carry out T trials, and in each trial, 70% of the training cases are randomly sampled, and used for training. A linear function $f_t(\vec{x}) = \vec{w}_t^T \vec{x} + b_t$ is then constructed in trial t . The final classifier is based on the averaged model $f(\vec{x}) = \frac{1}{T} \sum_{t=1}^T f_t(\vec{x})$. We remove features whose weights are very small in the function $f(\vec{x})$ in order to reduce the model noise. So, only the top N_f features with large weights remain in the final model. If a candidate \vec{x}_i achieves $f(\vec{x}_i) \geq a$, then it is classified as a true PE; otherwise, it is classified as a non-PE. a is the margin we set to specify the operating point on the ROC curve.

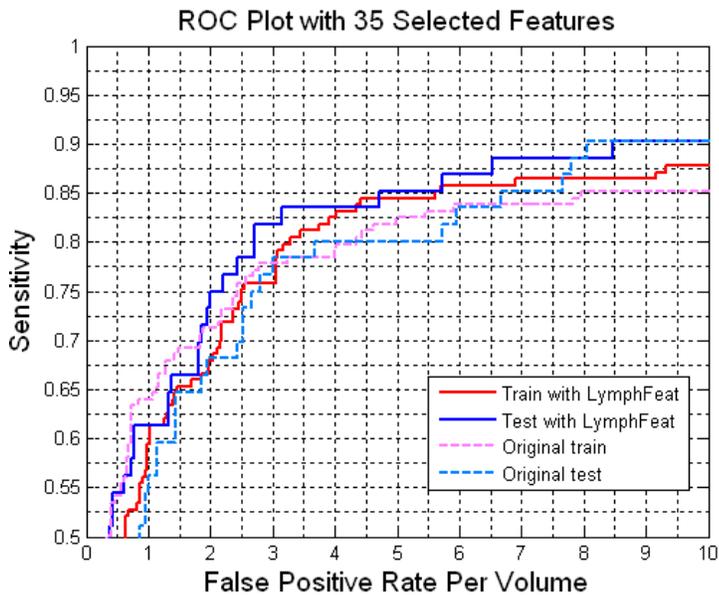
3. EXPERIMENTAL RESULTS

In this section, we demonstrate the effectiveness of the proposed lymph-specific features by performing two types of experiments. **(1)** We show that these features, when incorporated in the original detection system,¹ significantly reduce the FPs due to lymph candidates and improve the overall PE classification performance. **(2)** We show that these features are truly discriminative of lymph tissue when applied to lymph detection.

3.1 PE Detection and Lymph FP Reduction

Here, we compare the performance of our new system, which incorporates the the lymph-specific features, against our original system.¹ Figure 7 (a) plots the ROC results for the new (solid) and original (dashed) systems as applied to both training and testing samples. Here, we note that the top $N_f = 35$ most discriminative features are used to train and test both systems. We deduce that the new features lead to an overall sensitivity improvement of 3.4% at 4 FPs. Further, we perform an analysis of the performance at an ROC operating point of 78% sensitivity determined from the training cases. Figure 7(b) tabulates a numerical summary of the detection results for the training and testing cases. Under the original system, we provide the number of PEs detected, lymph FPs, and other FP types in the **Original** and **Common** columns, while under the new system, the results are given in the **Common** and **LymphFeat** columns. Note that the **Common** column shows the detections

common between the two systems. It is obvious that the number of lymph candidates that were incorrectly detected as PEs is significantly reduced, when the lymph-specific features are included. From these results, we also see that these features reduce the lymph FP rate by 59% while improving the overall sensitivity by 3.4% in the test cases. For the training cases, they reduce the lymph FP rate by 74%, while maintaining the same sensitivity as the original system.



TRAINING (45 cases, 2455 candidates)			
	Original	Common	LymphFeat
PE	6	115	7
Lymph FP	41	12	2
Other FP	37	94	42
Original FP rate = 4.09 per volume LymphFeat FP rate = 3.33 per volume			
TESTING (21 cases, 982 candidates)			
	Original	Common	LymphFeat
PE	1	46	3
Lymph FP	14	8	1
Other FP	21	37	21
Original Sensitivity = 78% Original FP rate = 3.81 per volume LymphFeat Sensitivity = 82% LymphFeat FP rate = 3.19 per volume			

Figure 7. (a) ROC curves for the original (dashed line) and new system (solid line) using lymph features. (b) performance analysis of both systems at an operating point of 78% sensitivity. We indicate the number of detected PEs, lymph FPs, and other FP types, as well as the number of common detections.

3.2 Lymph Detection

Here, we build a pure lymph FLD classifier, using the lymph-specific features ($N_f = 16$), whose performance further justifies the discriminatory power of these features. Figure 8 plots the ROC curves for both the training (45) and testing (21) cases. These curves show that with only these discriminative features lymph detection can be effectively performed.

4. CONCLUSION AND FUTURE WORK

In this paper, we present a new approach to reducing lymph false positives arising from the detection of pulmonary embolisms (PEs) in CT images of lungs. Novel contextual features were developed to describe the relative position of a candidate cluster of voxels to airways and arteries, which are anatomical features that primarily discriminate PEs from lymph. Embedding these new features into a multiple instance learning framework, we developed a machine learning algorithm to effectively classify PEs, while reducing FPs due to lymph. Our experimental results show the significant merit of these features in discriminating PE from lymph and in improving overall classification performance. In the future, we plan to extend this work to the detection of lymph tissue/nodules in CT images of parts of the body other than the lungs.

REFERENCES

1. J. Bi and J. Liang, "Multiple instance learning of pulmonary embolism detection with geodesic distance along vascular structure," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, 2007.

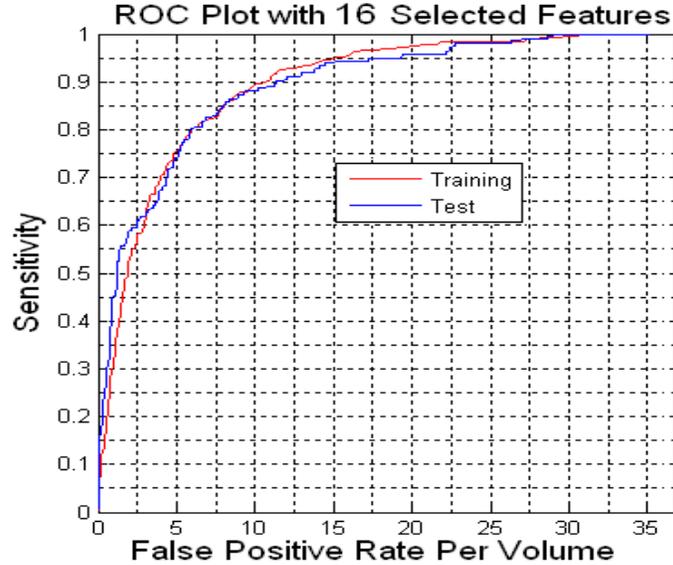


Figure 8. ROC curves for pure lymph classification.

2. Y. Masutani, H. MacMahon, and K. Doi, "Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis," *IEEE Transactions on Medical Imaging* **21**, pp. 1517–1523, 2002.
3. M. Quist, H. Bouma, C. van Kuijk, O. van Delden, and F. Gerritsen, "Computer aided detection of pulmonary embolism on multi-detector CT," in *Journals of the Radiological Society of North America*, 2004.
4. C. Zhou, L. M. Hadjiiski, S. Patel, H.-P. Chan, and B. Chan, "Computerized detection of pulmonary embolism in 3D computed tomographic (CT) images," in *RSNA*, p. (abstract), (Chicago, USA), Nov 2004.
5. S. Park, C. Bajaj, G. Gladish, and D. Cody, "Automatic pulmonary embolus detection and visualization," (poster at <http://www.ices.utexas.edu/~smpark>), 2004.
6. J. Liang and J. Bi, "Computer aided detection of pulmonary embolism with tobogganing and multiple instance classification in CT pulmonary angiography," *International Conference on Information Processing in Medical Imaging* **4584**, pp. 630–641, 2007.
7. J. Tschirren, E. A. Hoffman, G. McLennan, and M. Sonka, "Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose ct scans," *IEEE Transactions on Medical Imaging* **24**, pp. 1529–1539, 2005.
8. C. I. Fetita, F. Preteux, C. Beigelman-Aubry, and P. Grenier, "Pulmonary airways: 3-d reconstruction from multislice ct and clinical investigation," *IEEE Transactions on Medical Imaging* **23**, pp. 1353–1364, 2004.
9. C. Zhou, L. M. Hadjiiski, S. Patel, H.-P. Chan, and B. Chan, "Computerized detection of pulmonary embolism in 3D computed tomographic (CT) images," in *Journals of the Radiological Society of North America*, 2004.
10. M. P. Andrew W. Fitzgibbon and R. B. Fisher, "Direct least squares fitting of ellipses," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21**, pp. 476–480, 1999.
11. G. Medioni, M.-S. Lee, and C.-K. Tan, *A Computational Framework for Segmentation and Grouping*, Elsevier, 2000.
12. P. Mordohai and G. Medioni, *Tensor Voting: A Perceptual Organization Approach to Computer Vision And Machine Learning*, Synthesis Lectures on Image, Video, and Multimedia Processing, Morgan & Claypool, 2006.
13. G. Fung, M. Dundar, B. Krishnapuram, and B. Rao, "Multiple instance algorithms for CAD," *Neural Information Processing System*, 2006.