Computer Aided Detection of Pulmonary Embolism with Tobogganing and Mutiple Instance Classification in CT Pulmonary Angiography

Jianming Liang and Jinbo Bi

Computer Aided Diagnosis and Therapy
Siemens Medical Solutuions USA, Inc., Malvern, PA 19355

Abstract. Pulmonary embolism (PE) is a very serious condition causing sudden death in about one-third of the cases. Treatment with anti-clotting medications is highly effective but not without complications, while diagnosis has been missed in about 70% of the cases. A major clinical challenge, particularly in an Emergency Room, is to quickly and correctly diagnose patients with PE and then send them on to therapy. Computed tomographic pulmonary angiography (CTPA) has recently emerged as an accurate diagnostic tool for PE, but each CTPA study contains hundreds of CT slices. The accuracy and efficiency of interpreting such a large image data set is complicated by various PE look-alikes and also limited by human factors, such as attention span and eye fatigue. In response to this challenge, in this paper, we present a fast yet effective approach for computer aided detection of pulmonary embolism in CTPA. Our proposed approach is capable of detecting both acute and chronic pulmonary emboli with a distinguished feature of incrementally reporting any detection immediately once becoming available during searching, offering real-time support and achieving 80% sensitivity at 4 false positives. This superior performance is contributed to our novel algorithms (concentration oriented tobogganing and multiple instance classification) introduced in this paper for candidate detection and false positive reduction.

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

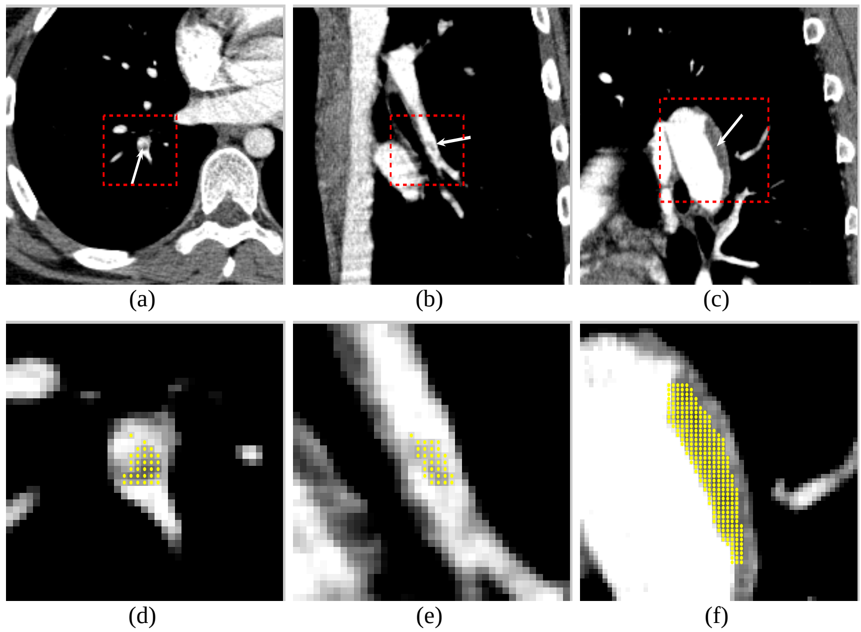

Fig. 1. The emboli appears as dark regions residing in bright vessel lumen. Our toboggan-based approach is capable to detect both acute (a, b) and chronic (c) pulmonary emboli, offering simultaneous detection and segmentation (d, e, f). The clot in (b) was actually missed by our radiologists, but correctly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

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 PEs 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) 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 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 [1–4]. 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 [1]. 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, large part of it would be excluded anyway. In this paper, we present a fast yet effective toboggan-based approach for automated PE detection in CTPA without vessel segmentation. Another distinguished feature of our approach is its highly interactiveness and real-time response. For clinical use, all the detections reported by a CAD systems must be reviewed and approved by radiologists. The existing PE CAD systems adopts a pipe-line architecture and only report the final detection at the end of execution, implying that the radiologist has to wait until the end of the system run in order to review and approve any detections. However, in an ER (Emergency Room) scenario, radiologists only have a small time window (2-3 minutes) to read a case and make the diagnosis. They cannot wait till the end of run to examine all the CAD detection. To meet this requirement, our approach is capable to report any PE detection once available in real time for radiologist to review and approve, while it continues searching for additional PEs. These capabilities are founded on our two novel algorithms introduced in this paper: concentration oriented tobogganing algorithm for candidate detection and mutiple instance classification algorithm for false positive reduction.

2 Basic tobogganing

Pulmonary embolism may be acute or chronic. They are only existing in pulmonary arteries and generally attached to the vessel wall (see Fig. 1). Due to the nature of their formation, CTPA reveals emboli, whether acute or chronic, as dark regions with Hounsfield Units (HU) between -50 HU and 100 HU, residing in contrast enhanced bright vessel lumen. However, due to partial volume effects, the pixels around the vessel boundaries are also in the same HU range. Therefore, a major challenge for automatic PE detection is to effectively separate the emboli from the vessel wall and to quickly remove partial volume effects around the vessel boundaries while correctly preserving the PE pixels. In response to this challenge, we have come up with an idea: sliding all the voxels with Hounsfield Units (HU) between -50 HU and 100 HU to its neighbor with minimal HU value and collecting all voxels that don't slide into regions with Hounsfield Unit below -50 HU. This idea is illustrated in Fig. 2 and explained in the following.

This algorithm is called tobogganing [5], which takes its name from the processing analogy of sliding down a steep hill and will be referred as "basic tobogganing" in this paper to be differentiated from a new tobogganing algorithm, called concentration oriented tobogganing, to be presented in Section 3. A basic operation in tobogganing is "sliding". A pixel v with intensity $P(v)$ and neighbors $N(v)$ slides down to pixel g:

$$g = \underset{\mathrm{t} \in \mathrm{N(v)} \cup \{\mathrm{v}\}}{\arg\min} \; P(t). \tag{1}$$

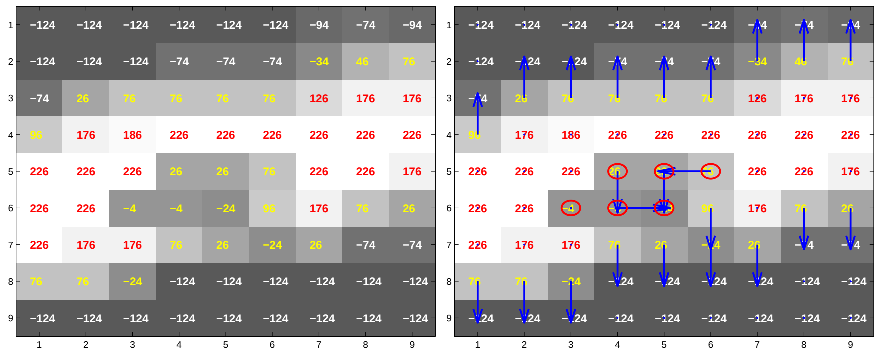

Fig. 2. An illustration of our idea for PE detection with basic tobogganing algorithm. In this small PE image, pixels with CT values below -50 HU are in white, pixels with CT value above 100 HU in red and all other pixels in yellow. Naturally, all the PE pixels are in yellow. However, due to partial volume effects, the pixels around the artery boundaries are also in yellow. Our idea can effectively detect the PE (circled) and remove the partial volume effects.

In cases where a pixel is surrounded by more than one pixel with the same minimal intensity value, the first pixel found with this value can be chosen or other more sophisticated strategies may be used in selecting a neighbor. A pixel that cannot slide to any of its neighbors is called a *concentration*. All the pixels that slide down to the same concentration form a toboggan *cluster* with a unique *label*.

The basic tobogganing algorithm operates as follows: Each unlabeled pixel slides till reaching a labeled pixel or a unlabeled new concentration. If it reaches a labeled pixel, the label is propagated back to all the pixels along the sliding path, otherwise, a new label is generated and then propagated back along the path. All the sliding directions may be recorded during the process. Referring to the simple PE image in Fig. 2, for detecting PE, we scan the image in row by row, but only selectively slide those pixels with CT values between -50 HU and 100 HU. For illustration, we use 2D four-connected neighborhood. The arrows indicate the sliding directions. During the tobogganing, the first pixel with CT value between -50 HU and 100 HU is pixel (7,2), which slides towards to pixel (7,1). Since the CT value of pixel (7,1) is -74 HU, and is pre-labeled as “air”, label “air” is propagated back to pixel (7,2). During the tobogganing process, the first pixel collected as PE pixel is pixel (4,5), because it slides down to pixel (4,6) and then concentrates at pixel (5,6) with CT value above -50 HU. Consequently, a PE label is generated and assigned to pixel (5,6) and propagated back to pixels (4,6) and (4,5). When it is done for all the pixels in the image, the yellow pixels around the arteries have all merged into darker regions with CT values below -50 HU and all the PE pixels stand out (circled in red). In this example, two toboggan clusters are formed for the detected PE pixels: The pixel (3,6) constitutes a single-pixel toboggan cluster, while all other pixels forms one cluster with concentration at pixel (5,6). To achieve the goal of PE detection, the adjacent toboggan clusters (with detected PE pixels) must be merged into a connected component, called a PE *candidate*, so that a detection position (3D point) can be derived by ultimate erosion to represent the candidate.

This basic tobogganing algorithm is intuitive and clearly useful in detecting PEs. However, a problem is that it only labels the PE voxels, providing suspicious PE regions. For PE detection, we must group the detected PE pixels into connected components, forming PE *candidates*. This means that we have to scan the whole 3D volumetric image data at least two times – one for tobogganing and one for connected component analysis, before reporting any detected PEs. In other words, the user (radiologist) has to wait the completion of two scans before reviewing and approving any PE detections. To overcome this drawback, we introduce concentration oriented tobogganing in the next section.

3 Concentration oriented tobogganing

3.1 The algorithm

During the PE search process, our goal is, once a PE pixel is encountered, to extract a whole PE candidate from the pixel immediately and send it to the user (radiologist) for review and approval. A PE candidate consists of multiple toboggan clusters. Naturally, in order to achieve the goal, we must first have an algorithm which can extract a toboggan cluster from any given pixel and provide the external boundary pixels of the cluster. The process of extracting a toboggan cluster from a given pixel is referred as concentration oriented tobogganning and formulated as an algorithm in Appendix which is iteratively invoked for detecting PEs. The idea is illustrated in Fig. 3 and detailed as follows.

Basically, the concentration oriented tobogganing algorithm has two steps. It first searches for concentration c from the given pixel s and then expands from the found concentration c to extract the whole toboggan cluster C. The expansion includes a base step and an iterative step. In the base step, it includes the concentration c as the first pixel in the cluster and pushes all its neighbors with CT values between -50 HU and 100 HU into an active list A. In the iterative step, it selects pixel q with the minimal CT value from the active list A, if the selected pixel toboggans to an already clustered pixel, then conditionally pushes its neighbors to the active list A to ensure the uniqueness of the pixels in the active list, otherwise, the selected pixel belongs to the cluster's external boundary B. The iterative step is repeated till the active list A is empty. This concentration oriented tobogganing algorithm is repeatedly applied on all those external boundary pixels, until a whole PE candidate has been extracted.

Referring to Fig. 3, when our example image is scanned in row by row, the first PE pixel encountered is (4,5), because it does not merge into regions with CT value below -50 HU. Therefore, we wish to extract the whole PE from the pixel at (4,5). Fig. 3.(a) illustrates Step A of the algorithm: finding the concentration. It regards the starting location as the current location, slides it to its neighbor with minimal CT value, then selects the neighbor as the current location and slides it until reaching a concentration. Once the concentration is found, if its CT value is between -50 HU and 100 HU, Step B is initiated to expand from the concentration to cover a whole toboggan cluster and provide all the external boundary pixels (circled in blue) as shown in Fig. 3.(b). In this example, the concentration (5,6) is included into the cluster and then all its neighbors with CT values in the PE HU range are pushed into an active list. A pixel with the

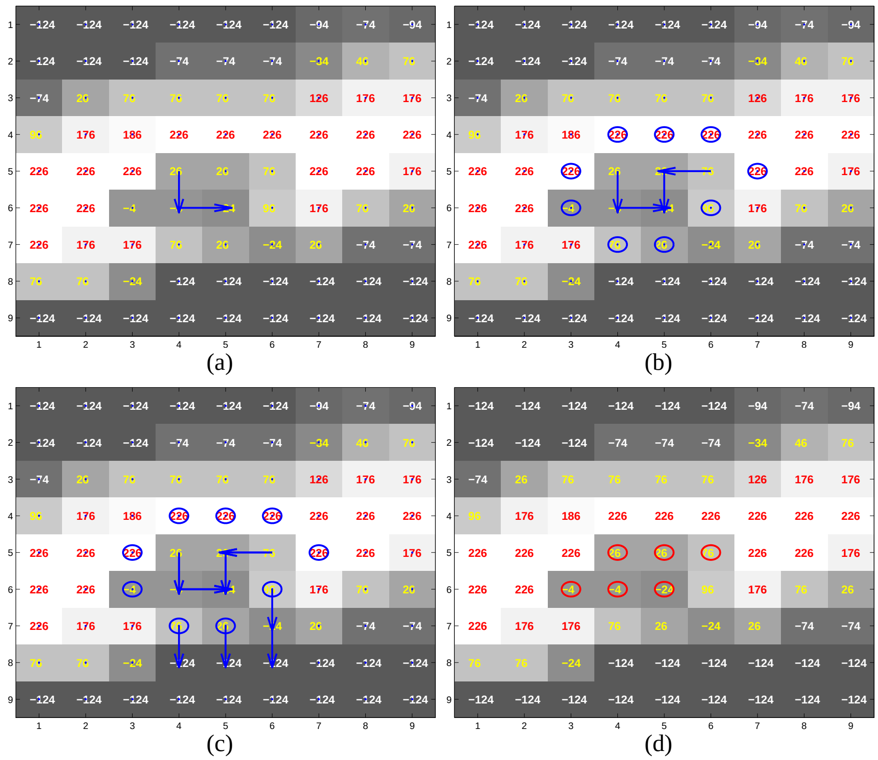

Fig. 3. Using the concentration oriented toboggan algorithm for PE detection. (a) Step A of the algorithm: Finding the concentration. (b) Step B: Expanding from the concentration to cover a whole toboggan cluster and providing all the external boundary pixels (circled). (c) Repeatedly apply the algorithm on all those external boundary pixels with CT value between -50 HU and 100 HU to form a PE candidate, leading to an identical result (d) as in Fig. 2.(b).

minimal CT value is selected from the active list. In this case, it is pixel (4,6). Since it slides towards pixel (5,6), which has been included into the cluster, its neighbors are conditionally pushed into the active list. The condition is that the neighbor must have CT value in the PE HU range, is not included in the cluster and is not in the active list. A new pixel with the minimal CT value is selected from the active list. For this time, it is pixel (3,6), but it does not slides towards any pixels within the cluster, therefore, it is a pixel on the external boundary of the cluster, and no processing is performed on its neighbors. Repeatedly select a new pixel from the active lust and process it in the same way till the active list is empty. Once done, we obtain all the pixels within the cluster, namely, (4,5), (4,6), (5,5), (5,6) and (6,5), as well as the pixels along the external boundary of the cluster (circled in blue in Fig. 3.(b)). The concentration oriented tobogganing algorithm is then iteratively applied on each of the external boundary pixels with CT value in the PE HU range. Any additional extracted toboggan cluster is merged into the previously extracted toboggan clusters, and any additional external boundary pixels are also merged. Once no external boundary pixel is left, all the toboggan clusters are extracted and merged, automatically forming a connected component – a PE candidate.

Proposition 31 *Concentration oriented tobogganing provides identical PE detections as basic tobogganing, but it has an advantage of reporting any detection immediately once becoming available during searching.*

3.2 Detection performance

We have collected 177 cases with 872 clots marked by expert chest radiologists at four different institutions (two North American sites and two European sites). They are divided into two sets: training (45 cases with 156 clots) and testing (132 cases with 716 clots). The training cases were used in the development process for algorithm developing, improving and testing, while the testing cases were only used for algorithm testing and were never used for development.

All the 177 cases were processed with our concentration oriented algorithm, which generated a total of 8806 candidates: 2431 candidates appear in the training set and 6375 candidates in the test set. Each candidate is a connected component – a cluster of voxels, and represented by a representative point with a 3-D coordinate derived from the cluster of voxels.

Each candidate was then labeled as a PE or not based on 3-D landmark ground truth provided by the experts. In order to automatically label each candidate, each PE pointed out by an expert landmark is semi-automatically extracted and segmented. Therefore, the ground truth for each PE is also a cluster of voxels (*i.e.*, the segmented PE). Any candidate that was found to be intersected with any of the segmented PEs in the ground truth was labeled as a PE. Multiple candidates may intersect with the same segmented PE, that is, multiple candidates may correspond to a single PE. Each PE is assigned with a unique identifier, therefore, multiple candidates may be labeled with the same PE identifier.

Our algorithms successfully detected 90.38% (141/156) of the PE in the training cases and 90.1%(645/716) of the PE in the testing cases. On average, the total computation time for each case is about 2 minutes on a 2.4GHz P4 PC and the first detection if any in a case is reported within 27 seconds. However, the concentration oriented algorithm also produces candidates that do not intersect with any PEs. These candidates are regarded as *false positives*. On average, 47.5 and 40.3 false positives for each case are generated for the training set and the test set, respectively. However, a system that "cries wolf" too often will be rejected out of hand by radiologists. Thus, the goal is to detect as many true PEs as possible, subject to a constraint on false positives, usually within 4 false positives per case. Therefore, we design a novel classification framework based on mathematical programming to reduce false positives in the next section.

4 False Positive Reduction

For clinical acceptability, it is critical to control false positive rates and detect as many true PEs as possible. A PE can be large, or have an elongated shape along the vessel, or split at the vessel bifurcation. Multiple candidate clusters may exist to intersect with a single PE. As long as one of the candidates is identified and visualized to physicians, the entire PE can be easily traced out. Consequently, it is sufficient to detect one candidate for each PE. Correct classification of every candidate instance is not as important as the ability to detect at least one candidate that points to a specific PE. Based on this concept, a novel multiple instance classification algorithm is devised to reduce false positives.

4.1 Feature Computation

A set of 116 descriptive properties, called *features*, are computed for each candidate. These features were all image-based features and were normalized to a unit range. The features can be categorized into several groups: those indicative of voxel intensity distributions within the candidate, those summarizing distributions in neighborhood of the candidate, and those that describe the 3-D shape of the candidate and enclosing structures. These features, in conjunction with each other, capture candidate properties that can disambiguate true emboli from typical false positives, such as dark areas that result from poor mixing of bright contrast agents with blood in veins, and dark connective tissues between vessels. These features are not necessarily independent, and may be correlated with each other, especially within the same group.

4.2 Mutiple Instance Classification

Assume that totally ℓ candidates are extracted, each represented by a feature vector $\mathbf{x}_i$ associated with a label y_i. The label $y_i = 1$ if the candidate overlays on a PE, or otherwise $y_i = -1$. Let I^+ and I^- be two index sets containing indices of candidates that intersect with PEs and do not intersect with PEs, respectively. Let m be the total number of PEs marked by expert radiologists for the n images. Denote I_j as the index set of the candidates that intersect with the j-th PE, $j = 1, \cdots, m$. Notice that $\cup_{j=1,\cdots,m} I_j = I^+$ but any two index sets I_j's are not necessarily disjoint since there may exist a candidate cluster that intersects with more than one segmented PEs.

Support vector machine (SVM) [6–8] has been a successful methodology for classification and regression. It constructs linear classification functions of the form $\mathbf{w}^T\mathbf{x} + b$ by minimizing the hinge error defined as $\xi = \max\{0, 1 - y(\mathbf{w}^T\mathbf{x} - b)\}$ for all candidates. We derive a more effective classification approach by exploring the key observation that once a candidate in I_j is classified as a positive, then the j-th PE is considered being identified. This consideration suggests the classifier to focus on different PEs instead of multiple candidates within a single PE. Especially it facilitates the reduction of false positives by possibly ignoring extremely noisy candidates that intersect with some PEs where, for the same PE, other associated candidates can be easily classified correctly. A geometric interpretation is illustrated in a 2-D feature space as in Fig.4 where standard SVMs focus on separating all candidates to correct sides whereas our learning algorithm classifies at least one true PE candidate into one side and others on the other side, thus successfully removing all false detections.

Mathematically, distinguishing at least one candidate for each PE from the negative class is equivalent to the statement that as long as the minimum of the errors (ξ) that are occurred on the candidates associated with a PE is 0, then that PE is detected. For

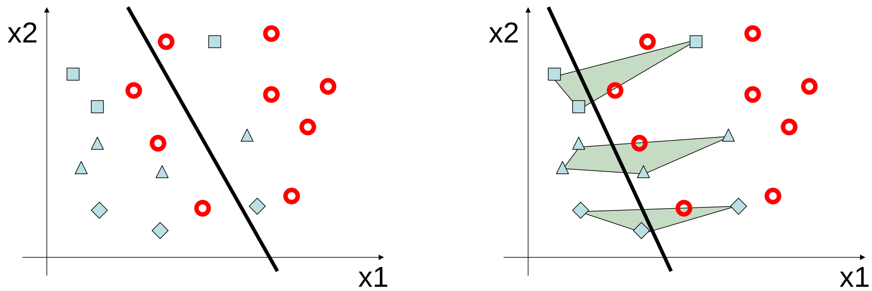

Fig. 4. Illustration of the classification. (Left) the linear separation boundary by standard SVM where circles represent false detections, and the symbols (diamond, box and triangle) each represent one PE with multiple candidates. (Right) the linear separation boundary by our multiple instance classification algorithm with more significant false positive reduction.

example, if a PE is associated with 3 candidates, and a classifier generates $\xi_1 = 0$ for the first candidate, $\xi_2 = 5,\ \xi_3 = 100$ for the other two candidates, the classifier detects the PE. Correspondingly, this implies to construct the classifier by solving the following optimization problem:

$$\begin{array}{ll} \min_{\mathbf{w},\xi} & \gamma||\mathbf{w}||_1 + \sum_{j=1}^{m} \min\{\xi_i,\ i \in I_j\} + \sum_{i\in I^-} \xi_i \\ \text{s.t.} & \mathbf{w}^T x_i + b \geq 1 - \xi_i,\ i \in I^+, \\ & \mathbf{w}^T x_i + b \leq -1 + \xi_i,\ i \in I^-, \\ & \xi_i \geq 0,\ i = 1, \cdots, \ell. \end{array} \tag{2}$$

However, this optimization problem is computationally difficult to solve, because it involves a minimization of the to-be-determined variables ξ in the evaluation of the objective function, and it is neither differentiable nor convex. Hence, it is necessary to devise a tractable optimization problem that is equivalent. To this end, we prove that problem (2) is equivalent to the quadratic program (3), as characterized by the following theorem:

$$\begin{array}{ll} \min_{\mathbf{w},\xi,\boldsymbol{\lambda}} & \gamma||\mathbf{w}||_1 + \sum_{j=1}^{m} (\sum_{i\in I_j} \lambda_i \xi_i) + \sum_{i\in I^-} \xi_i \\ \text{s.t.} & \mathbf{w}^T x_i + b \geq 1 - \xi_i,\ i \in I^+, \\ & \mathbf{w}^T x_i + b \leq -1 + \xi_i,\ i \in I^-, \\ & \xi_i \geq 0,\ i = 1, \cdots, \ell, \\ & \sum_{i\in I_j} \lambda_i = 1, \lambda_i \geq 0,\ i \in I_j,\ j = 1, \cdots, m. \end{array} \tag{3}$$

Theorem 41 *An optimal solution* $\hat{\mathbf{w}}$ *of Problem (2) is also optimal to Problem (3) with properly chosen* λ*, and vice versa.*

Proof. First of all, we prove that an optimal solution of Problem (3) has nonzero λ's only on the candidates for which the classifier achieves $\min\{\xi_i,\ i \in I_j\},\ \forall\, j$.

Let $(\hat{\mathbf{w}}, \hat{\boldsymbol{\xi}}, \hat{\boldsymbol{\lambda}})$ be the optimal solution of Problem (3). For notational convenience, denote the objective of Problem (3) as $\mathcal{J}(\mathbf{w}, \boldsymbol{\xi}, \boldsymbol{\lambda}) = \gamma||\mathbf{w}||_1 + \sum_{j=1}^{m}(\sum_{i\in I_j} \lambda_i \xi_i) + \sum_{i\in I^-} \xi_i$. Then let $\hat{\mathcal{J}}$ be the objective value attained at the optimal solution. Notice

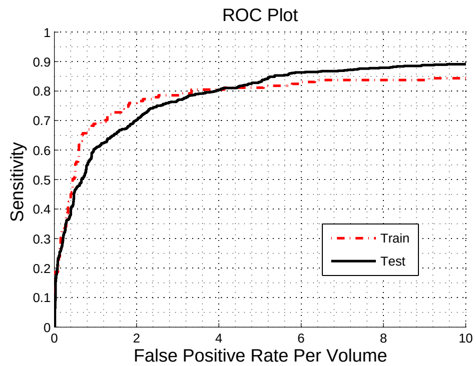

Fig. 5. The ROC plot of the final system.

that the hinge loss $\hat{\boldsymbol{\xi}}$ is uniquely determined by $\hat{\mathbf{w}}$ as $\hat{\xi}_i = \max\{0, 1 - y_i(\mathbf{w}^T\mathbf{x}_i + b)\}$ for each candidate $\mathbf{x}_i$.

If $\exists\, j \in \{1, \cdots, m\}$, and $\exists i_0 \in I_j$, such that $\lambda_{i_0} > 0$ but $\hat{\xi}_{i_0} \neq \min\{\xi_i, i \in I_j\}$. Then let $\xi_{I_j} = \min\{\xi_i, i \in I_j\} < \hat{\xi}_{i_0}$. Then $\tilde{\mathcal{J}} = \hat{\mathcal{J}} - \lambda_{i_0}\xi_{i_0} + \lambda_{i_0}\xi_{I_j} < \hat{\mathcal{J}}$. This contradicts to the optimality of $(\hat{\mathbf{w}}, \hat{\boldsymbol{\xi}}, \hat{\boldsymbol{\lambda}})$.

By this contradiction, $\forall i, j$, such as $\lambda_i > 0$, the corresponding ξ_i has to be the minimum loss that the classifier achieves on the j-th PE. This implies that at the optimality of Problem (3), $\mathcal{J} = \gamma||\mathbf{w}||_1 + \sum_{j=1}^{m} \min\{\xi_i,\ i \in I_j\} + \sum_{i \in I^-} \xi_i$ which is the same as the objective of Problem (2).

4.3 System Performance

Our classification algorithm has dramatically reduced the false positive rate down to 4 false positives per patient while maintaining the high detection sensitivity. Fig.5 depicts the Receiver Operating Characteristics (ROC) plot of our final system that combines the candidate detection, feature computation, and classification. As shown in Fig.5, the final system detects 80% of the PEs, respectively, for the training study set and the test set at 4 false positive per patient.

5 Discussions and Conclusions

We view our method as toboggan-based because the idea is originally inspired by the work of Fairfield [5] and of Mortensen and Barrett [9]. The basic tobogganing algorithm presented in Section 2 is a modified version of the algorithm in [5, 10]. Nevertheless, each of the research groups has different aims in mind. Fairfield aimed to enhance the contrast of images by tobogganing, while Mortensen and Barrett used tobogganing with the aim to group the pixels to reduce the underlying graph in livewire for efficient interactive image segmentation. Clearly, our aim is to separate objects (PEs in this case) from adjacent (connected) objects (vessel walls in this case) and remove partial volume effects. Given the general nature of the idea, the algorithm has been successfully applied to other applications, for instance, detecting colonic polyps in CT images.

We also would like to contrast our concentration oriented tobogganing algorithm with a few of existing related algorithms in the literature including: watershed, hierarchical tobogganing, intelligent paint, and intelligent scissor" (*i.e.*, "live-wire"). There is a rich set of algorithms in the watershed literature. The most related ones are rainfalling simulation [11] and the watershed technique based on hill climbing reported in [12]. The basic toboggan algorithm first reported by Fairfield largely went unnoticed in the watershed community. Rainfalling simulation can be regarded as an extension of Fairfield's algorithm for handling "plateau". The watershed technique based on hill climbing reported in [12] requires that *all* the minima be found in advance and marked with distinct labels followed by "hill climbing". This implies that we would not be able to obtain a watershed region till the whole image has been scanned and processed. "Hierarchical tobogganing" is to repeatedly apply the basic toboggan algorithm, forming toboggan hierarchy. "Intelligent paint" is built on top of hierarchical tobogganing to allow the user to interactively "select" the pre-formed toboggan regions at a user pre-specified toboggan hierarchical level, based on cost-ordered region collection. "Intelligent scissor" or interactive "live-wire" aims to compute an optimal path from a selected seed point to *every* other point in the image based on unrestricted graph search, so that the user can move the mouse freely in the image plane and interactively "select" a desired path among all the optimal paths based on the current cursor position. The underlying algorithm is Dijkstra's algorithm, which computed a shortest path from a given point to every other point in the image. However, for large images, the underlying graph created in live-wire for search become large, the interactiveness of livewire would be comprised due to the fundamental limitation of Dijkstra's algorithm. Therefore, Mortensen and Barrett [9] proposed toboggan-based livewire, in which the basic toboggan algorithm is applied to reduce the underlying graph in livewire to achieve highly efficient interaction in image segmentation. In short, all the discussed algorithms cannot meet our requirement to extract a toboggan cluster from an initial site without processing any pixels beyond its external boundary.

In conclusion, we have developed a novel approach for computer aided detection of pulmonary embolism. Our approach has a set of distinguished features, requiring no vessel segmentation, reporting any detection incrementally in real time, and detecting both acute and chronic pulmonary emboli, achieving 80% sensitivity at 4 false positives.

References

1. 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*, vol. 21, no. 12, pp. 1517–1523, 2002.
2. M. Quist, H. Bouma, C. van Kuijk, O. van Delden, and F. Gerritsen, "Computer aided detection of pulmonary embolism on multi-detector CT," in *RSNA*, Chicago, USA, Nov 2004, p. (Conference Poster).
3. 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*, Chicago, USA, Nov 2004, p. (abstract).
4. S. Park, C. Bajaj, G. Gladish, and D. Cody, "Automatic pulmonary embolos detection and visualization," poster at http://www.ices.utexas.edu/∼smpark, 2004.

5. J. Fairfield, "Toboggan contrast enhancement for contrast segmentation," in *IEEE Proceedings of 10th International Conference on Pattern Recognition*, vol. 1, Atlantic City, June 1990, pp. 712–716.
6. V. Vapnik, *The Nature of Statistical Learning Theory*. New York: Springer, 1995.
7. B. E. Boser, I. M. Guyon, and V. N. Vapnik, "A training algorithm for optimal margin classifiers," in *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory*, D. Haussler, Ed., 1992, pp. 144–152.
8. J. Zhu, S. Rosset, T. Hastie, and R. Tibshirani, "1-norm support vector machines," in *Advances in Neural Information Processing Systems 16*, S. Thrun, L. Saul, and B. Schölkopf, Eds. Cambridge, MA: MIT Press, 2004.
9. E. N. Mortensen and W. A. Barrett, "Toboggan-based intelligent scissors with a four parameter edge model," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Fort Collins, CO, June 1999, pp. 452–458.
10. E. N. Mortensen, "Simultaneous multi-frame subpixel boundary definition using toboggan-based intelligent scissors for image and movie editing," Ph.D. dissertation, Department of Computer Science, Brigham Young University, Provo, UT, 2000.
11. S. Stoev, "Rafsi - a fast watershed algorithm based on rainfalling simulation," 2000. [Online]. Available: citeseer.ist.psu.edu/stoev00rafsi.html
12. J. Roerdink and A. Meijster, "The watershed transform: Definitions, algorithms and parallelization techniques," 1999. [Online]. Available: citeseer.ist.psu.edu/roerdink01watershed.html

Appendix: A concentration oriented tobogganing algorithm

```
Input:
   s                 {initial site}
   p = P(v)          {toboggan potential p of voxel v}
   lt                {the low threshold}
   ht                {the high threshold}
Output:
   C                 {toboggan cluster containing initial site s; initially empty}
   B                 {external boundary pixels of cluster C; initially empty}
Data Structures:
   A                 {Active list of voxels; initially empty}
Functions:
   E = N(v)          {get neighbors E of voxel v}
   g = tob(v)        {g = arg min_{t ∈ N(v) ∪ {v}} P(t) }
   A = update(A,v)   {∀ r ∈ N(v), A ← r, if (r ∉ A) and (r ∉ C) and (P(r) ∈ [lt, ht]) }
   q = pop(A)        {q = arg min_{a∈A} P(a) }
Steps:
   {Step A: Find concentration c of initial site s}
      c = s;
      repeat
         q = c; c = tob(q);
      until (q = c)
   {Step B: Expand from concentration c}
      if (P(c) ∈ [lt, ht]) begin
         {Step B.1: Base step}
            C ← c; A = update(A, c);
         {Step B.2: Iterative step}
            repeat
               q = pop(A); r = tob(q);
               if r ∈ C begin {include q into cluster C and update A}
                  C ← q; A = update(A,q);
               end else begin {include q into external boundary B}
                  B ← q;
               end
            until A is empty
      end
```