Published in final edited form as:

Acad Radiol. 2009 February ; 16(2): 172–180. doi:10.1016/j.acra.2008.07.009.

Computer Aided Detection of Endobronchial Valves using Volumetric CT

Robert A. Ochs, MS,

Department of Radiological Sciences, David Geffen School of Medicine, UCLA, Address: 924 Westwood Blvd., Suite 650, Los Angeles, CA 90024-2926, Telephone number: (310) 794-8968, Fax Number: (310) 794-8657, Email address: rochs@mednet.ucla.edu

Fereidoun Abtin, MD,

Department of Radiological Sciences, UCLA

Raffi Ghurabi, MS,

Department of Radiological Sciences, UCLA

Ajary Rao, BS,

Department of Radiological Sciences, UCLA

Shama Ahmad, MS,

Department of Radiological Sciences, UCLA

Matthew S Brown, PhD, and

Department of Radiological Sciences, UCLA

Jonathan G Goldin, MbChB, PhD

Department of Radiological Sciences, UCLA

Abstract

Rationale and Objectives—The ability to automatically detect and monitor implanted devices may serve an important role in patient care by aiding the evaluation of device and treatment efficacy. The purpose of this research was to develop a system for the automated detection of one-way endobronchial valves that were implanted for less-invasive lung volume reduction.

Materials and Methods—Volumetric thin-section CT data was obtained for 194 subjects; 95 subjects implanted with 246 devices were used for system development and 99 subjects implanted with 354 devices were reserved for testing. The detection process consisted of pre-processing, pattern-recognition based detection, and a final device selection. Following the pre-processing, a set of classifiers were trained using AdaBoost to discriminate true devices from false positives. The classifiers in the cascade used two simple features (either the mean or maximum attenuation) of a local region computed at multiple fixed landmarks relative to a template model of the valve.

Results—FROC analysis was performed for the evaluation; the system could be set so the mean sensitivity was 96.5% with a mean of 0.18 false positives per subject. If knowledge of the number of implanted devices were incorporated, the sensitivity would be 96.9% with a mean of 0.061 false positives per subject; this corresponds to a total of 12 false negatives and 6 false positives for the 99 subjects in the test dataset.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conclusion—Software was developed for automated detection of endobronchial valves on volumetric CT. The proposed device modeling and detection techniques may be applicable to other devices as well as useful for evaluation of treatment response.

Keywords

Computer Aided Detection; Endobronchial Valves; CT; AdaBoost Cascade

1. INTRODUCTION

Initial studies of a one-way endobronchial valve (Emphasys Medical Inc., Redwood City, CA) for less-invasive emphysema treatment have recently been completed with favorable results, although the results from a large multicenter clinical trial are still pending [1–3]. Previous approaches for emphysema treatment have been palliative therapy or lung volume reduction surgery. The objective of the new endobronchial valve therapy, and other newly emerging alternative therapies, is to place multiple one-way valves into the airways of a targeted, diseased lobe of the lung in order to prevent entry of air into the lobe, while still allowing air inside the lobe to flow out [4]. By only allowing air out, the treatment aims to reduce the volume of the lobe so that neighboring, healthier lung tissue will expand, providing a treatment benefit to the patient.

The motivation for this research is to improve the care of patients implanted with medical devices. The ability to detect and monitor implanted devices is important for the evaluation of device and treatment efficacy. Treatment plans may provide guidance to physicians of possible treatment locations; however, either intentionally or unintentionally, physicians may not always place devices in the planned locations. Additionally, once implanted, devices may not always be stable; they can migrate into surrounding lung parenchyma or be coughed out. Figure 1 shows an example of a valve that migrated into the parenchyma, as well as another example of a valve that was placed too distal for complete occlusion of the lobe.

Computer-aided detection and monitoring of implanted devices from volumetric CT data is an emerging field. Techniques have been developed for detecting fiducial markers on 2D projection images to aid target location for radiation therapy [5]. A template matching technique has been developed to detect paramagnetic markers on 2D MR projection images [6]. Markers are often comprised of solid high attenuation material, allowing them to be detected on projection images; however, endobronchial valves are comprised of thin wires, more complex geometry, and can be obscured by local tissue or calcifications, all of which can increase the difficulty of the detection task.

The primary aim of this research is to develop a technique to automatically detect endobronchial valves. The underlying methods utilize generic techniques, which may potentially be adapted to other devices as well.

2. METHODS

2.1 Device detection

The detection process consisted of three main steps: 1) pre-processing, 2) detection, and 3) device selection. While the pre-processing is likely to be domain specific, the underlying detection algorithm is intended to be generic enough for many devices. The detection process passes a template of the device at multiple orientations through each test point remaining after the pre-processing. Features computed at fixed landmarks (i.e., *control-points*) on each template are used for differentiation of true devices from false positives.

2.1.1 Domain specific pre-processing—The objective of the pre-processing step is to efficiently reduce the number of points that need to be tested during the detection process. Each CT image contains 512×512 pixels, and if there are 512 axial images, there would be over 134 million points to test with computationally intensive 3D detection techniques. The goal is to have less than 10,000 points, making it feasible for the detection process to occur in a matter of seconds.

Segmentation of the appropriate anatomical region and application of a threshold can significantly reduce the number of test points. Using a previous developed thoracic segmentation algorithm, the spine is subtracted from the lung-hilum region [7]. A threshold is then applied; voxels with an attenuation value above the threshold became part of the initial set of test points.

The appropriate threshold was determined empirically using the known attenuation and samples in the developing dataset. It is desirable to have the highest threshold as possible to reduce the amount of noise in the processing and/or false positives in later stages of the device detection process. For each sample in our training data, the threshold was incrementally raised until the device was not longer part of the set of test points. For images of 0.6 mm slice thickness, all devices could be found using a threshold of 600 HU. For images of 1.5 mm slice thickness, the threshold needed to be lowered to 250 HU due to partial volume averaging of the device.

2.1.2 Control points and features—A control point model of the device was created to represent a physical model of the endobronchial valve. The valve was modeled by a set of n control points, $\{\vec{p}_0 \dots \vec{p}_n\}$, where $\vec{p}_i$ denotes the spatial location (x_i, y_i, z_i) of each control point relative to an origin point on the device. A set of templates were formed by rotating the model to unique orientations; for the valves, each template was orientated to one of the 98 surface vertices of a $5 \times 5 \times 5$ cube surrounding a point were used. The term *point-template* refers to a unique orientation of the template at a given test point remaining from the pre-processing.

The physical device model was determined using valves from the training data. The endobronchial valve was modeled simply as a small base cylinder connected to a larger outer cylinder. Electronic calipers were used to manually measure the dimensions of the base and outer cylinders on training samples. A 3D model of the device based on the manual measurements was created and overlaid onto valves in the training dataset for visually confirmation that the model matched the CT attenuation data. Even spaced control points were placed on the base and outer cylinders of the model (i.e., areas of higher attenuation). Additional points were placed in the lumen of the device (i.e., areas of lower attenuation).

Each control point was assigned a group number (i.e., to a control-group). Group assignments were based on where the point was located on the physical model of the valve; points in the same group are expected to have very similar features. For instance, as seen in the illustration the device model in Figure 2, points on the base of the device are in group 1, points on the outer cylinder are in group 2, and points in the lumen of the device are in group 3.

The primary reason for using control groups was to increase the number of training samples relative to the dimensionality of the training vector. For instance, if there were 58 control points, and two possible attenuation features at each control point, it would result in 116 dimensions for training; however, if the training is done by group, and there are 4 control groups with 2 features per group, then there would be only 8 training dimensions. In addition, by combining control points, each individual control point sample becomes part of the aggregate set of training samples for its respective group, which results in a multiple of (n -control-points over

n-control-groups) times the number of samples for training. For the above example, this results in approximately 15 (e.g., 58/4) times more samples per training dimension.

For this study, two basic attenuation features were computed at each control point for classification: the maximum attenuation, $I_{\max}(\bar{p}_i)$, and mean attenuation, $I_{\text{avg}}(\bar{p}_i)$, over the resolution of the control point. The resolution refers to the number of voxels surrounding the control point being tested; the resolution is the same for all points in a control group. When implanted, devices may have some slight deformations or may not align exactly with one of the templates; the resolution is used to allow for this slight variance from the underlying device model. Appropriate control points and resolutions, either $3 \times 3 \times 3$ or $5 \times 5 \times 3$, were determined by overlaying the device model onto training samples and visually confirming their alignment and coverage.

2.1.3 AdaBoost machine learning—The AdaBoost machine learning algorithm was used to train a cascade classifier for detecting endobronchial valves [8–10]. The AdaBoost algorithm was selected based on its strong theoretical basis, simplicity to implement, transparency of feature selection, and performance for other discriminative tasks [11]. The algorithm concurrently selects and combines relevant features from the feature set during the training each classifier, thus avoiding a separate feature selection process common with other classification methods. The basic premise of the algorithm is that any number of weak classifiers with an error rate less than 50% can be combined to form an ensemble classifier whose error rate approaches that of an optimal classifier.

For the valve detection, the set of weak classifiers consisted of one weak classifier per control point feature. An exhaustive search was used to train the weak classifiers; the optimal threshold condition of the weak classifier was determined within the range of -1020 HU to 1020 HU for both the maximum and mean features.

A slight modification of the algorithm was made to fit the device detection process. Normally, each weak classifier is trained independently; however, as previously mentioned, control points are assigned to a control group, and modifications of the underlying training code were made so that the algorithm would train weak classifiers using samples from all points in a control group. This results in control points sharing the same classification decisions as all other points in their respective control group during each iteration.

The AdaBoost cascade was comprised of four levels; the number of features in each level of the cascade was as follows: 4, 8, 16, and 32. During implementation, an instance is tested with the first level; if negatively classified, the instance is rejected; if positively classified, the instance is tested with the next level, and so forth until rejected or it has passed through all of the cascade levels. The control point feature and cascade approach are very efficient; for example, for the first cascade only 4 features need to be computed, if positively classified, then the next 8 features are computed, etc... As discussed next, the vast majority of instances can be rejected with just a few simple operations in the first cascade, thus avoiding the computation of features at all control points for classification.

The decision threshold of the classifier output for each level of the cascade was set just low enough so there would be no false negatives during training. Initially, the set of true negative training samples was comprised of all remaining points following the preprocessing from the training series (minus the true devices); this resulted in over 100,000 negative points. Since each point can have 98 possible template instances, this corresponds to over 11,000,000 negative instances. For training each level of the cascade, all positives were used and approximately 100,000 negative instances (or as many that remained) were randomly selected from the pool of negatives. After training a level, the classifier was then re-applied to the

training samples; instances positively classified by the cascade became part of the pool of negatives that were randomly sampled for training the next level of the cascade.

Table 1 and Figure 3 illustrate the first four weak classifiers and their underlying control point features, selected by the AdaBoost algorithm for the first level of the cascade classifier. The top features were the maximum attenuation over the local control point resolution. Control points on the outer ring and inner lumen were likely selected because they help differentiate the valve structure from calcifications.

The efficiency of the cascade approach comes from the fact that 90% of negative instances can be rejected with the first level, and 99% of negatives can be rejected after the second level. Table 2 reports the system performance during training; each additional level of the cascade caused an approximate 10-fold reduction in negative instances. By the fourth cascade, there were only 1197 false positives from over 11,000,000 initial negative examples. To avoid over fitting, it was decided that only four cascades should be used. The device confidence score, discussed next, was used to select the final set of devices.

2.1.4 Device confidence score—Since the valves are often grouped together, a final generic device confidence score, σ , was used to separate and identify individual valves. The confidence score is only computed for instances that positively passed through all levels of the AdaBoost cascade. A simple scoring scheme that used all control points was found to be better for computing the final device score than the relatively smaller subset of control points used by the AdaBoost cascade for detection. The algorithm for computing the confidence score is shown in Figure 4. The confidence score formed from the summation of individual scores based on the maximal attenuation at each control point, $I_{\max}(\bar{p}_i)$. For instance, in Figure 2, groups 1 and 2 were modeled as on the edge of the device, while control points in group 3 are in the device lumen (i.e., not on the metallic framework). The values of the scoring were determined heuristically using the training samples.

2.2 Implementation

During implementation, the image data is loaded, segmented, and the set of test points are generated. The detection algorithm presented in Figure 5 is then implemented to generate the final set of valves. An illustration of the process is shown in Figure 6. At each test point, each of the 98 template instances is tested with the AdaBoost cascade classifier. Confidence scores are computed only for instances that positively passed through all levels of the cascade. A list of each device with its confidence score is created. The final set of devices is formed by accepting devices from the list in descending order according to their confidence score; however, devices that overlap an existing device are rejected. Domain knowledge can also be used during this process; for instance, if known there should be only d implanted devices then only the highest d non-overlapping devices can be accepted from the list.

2.3 Endobronchial valve dataset

The dataset was composed of volumetric CT data from 194 subjects scanned at inspiration. Subjects were imaged on one of ten scanner models from four commercial vendors: General Electric (GE Medical Systems; Waukesha, WI); Philips (Philips Healthcare, Andover, MA); Siemens (Siemens Medical Solutions, Forchheim, Germany); and Toshiba (Toshiba America Medical Systems, Tustin, CA). Image acquisition protocols were standardized across all models to maintain comparable image quality; scanning centers were instructed to use a standardized scan protocol of 120 kVp, 150 to 300 mA, and a pitch of 1.375 to 1.5. Images were reconstructed with a display field of view containing the widest cross-section of the lungs, using non-enhancing reconstruction algorithms (e.g., GE Bone algorithm, Philips D, Siemens B50, or Toshiba FC51). The slice thickness of the image series was dependent on the capability

of each scanner model and was either 0.6 mm (n=17), 1.0 mm (n=16), 1.25 mm (n=94), or 1.5 mm (n=67). The slice interval was either < 1.0 mm (n=163), 1.0 mm (n=23), or > 1.0 mm (n=8).

As part of a clinical trial, subjects were implanted with the Zephyr® one-way endobronchial valves (Emphasys Medical Inc., Redwood City, CA). The one-way valve consists of a metal, diamond shaped framework embedded in a silicone retainer for stabilizing the device in the airways; the dimensions of the valves are approximately 8–10 mm in diameter and 15 mm in length. Multiple valves are implanted in the segmental airways of a targeted treatment lobe; in our dataset, the number of valves ranged from 1–9 per subject with a median of 3.

The subjects were randomly separated into a development dataset and testing dataset (36 subjects were automatically placed into the development dataset because they had been utilized for a previous pilot study). The development dataset consisted of 95 subjects total, of whom, 69 subjects had been implanted with a total of 246 valves. Subjects without implanted valves were used for training the system against false positives. The test dataset consisted of 99 subjects who had been implanted with a total of 354 valves. The dataset were kept separate; no testing subjects were used during development and no development subjects were used for testing.

2.4 Statistical analysis

The standard of truth for the number of implanted valves came from radiologist reads of each series, in conjunction with a technologist (medical student) review of the system. A radiologist read each case without the computer-aided device detection and recorded the number of implanted valves. Two medical students then reviewed the system output, with knowledge of the number of valves found by the radiologists. The medical students opened the image data, reviewed the computer output, and identified correctly marked devices (true positives), incorrectly marked devices, (false positives), and missed devices (false negatives). The results are reported for the initial evaluation of the system on the test dataset. Results of the system are shown with a FROC curve; the confidence score of the device served as the decision value for the FROC analysis.

3. RESULTS

The FROC for the device detection performance is shown in Figure 7. Two curves are shown (a) and (b); each curve shows the relationship between the sensitivity of the system and the false positive rate of the system. For example, when the system has a high sensitivity, it will find more devices, but the false positive rate will be higher. Conversely, the computer can be made less sensitive, thus lowering false positive rate and possibly making the system output easier to review. Curve (a) on Figure 7 shows that the system could be set so the mean sensitivity was 96.5% (338 of 354 correctly detected devices) with a mean of 0.18 false positives per subject; these results are when the system does not know how many devices were implanted.

Curve (b) on Figure 7 shows that the performance of the system can be improved when it has knowledge of the number of implanted devices. It can be expected that physicians should record the number of valves they have implanted. The computer can then return a maximum of the number of implanted valves during the detection process; for instance, if the physician implanted five valves, then logically, the computer should return no more than five valves during the detection process. By incorporating this prior knowledge, the sensitivity essentially remains the same 96.9% (342 of 354 correctly detected devices), but the false positive rate decreases to a mean of 0.061 false positives per subject. For the 99 subjects in the test dataset, this corresponded to a total of 12 false negatives and 6 false positives.

Figures 8–10 show examples of true positives, false positives, as well as false negatives. Figure 8 shows examples of two correctly detected devices in collapsed lung. Figure 9 shows example of two missed device anterior to two found devices. The slice thickness and orientation of the devices may have resulted partial volume averaging, resulting in the devices being missed by the initial threshold. Figure 10 shows examples of false positives due to metal artifacts from pacemaker leads.

4. DISCUSSION

The Food and Drug Administration (FDA) has recognized the role imaging can play for monitoring safety and efficacy of implanted devices. The development of practice guidelines for monitoring implanted devices was one of seventy-six research goals listed by the FDA in its 2006 Critical Path Opportunities List for the improvement of healthcare [12]. Other studies have confirmed the potential for imaging to offer a less invasive approach for device monitoring. This is exemplified by one study where one month follow-up bronchoscopic exams were initially used to evaluate device locations; however, due to poor compliance with the procedure, it was replaced with non-invasive CT imaging [3]. A recent study has also demonstrated the potential of high-resolution CT for the evaluation of coronary stents as opposed to conventional coronary angiography studies [13].

Initial analysis of the efficacy of the endobronchial valve therapy has shown the treatment has been effective for some subjects; however, not all subjects have had a complete lobar collapse [1–3]. Understanding how the number and placement of devices influence the treatment outcome may be beneficial for both improving the treatment planning process, as well as for developing guidelines as to when follow-up action is required due to incorrect device placement.

The results of our study indicate a strong potential for computer systems to play a key role in automating the device monitoring process for implanted endobronchial valves. This research may be applicable to other devices that are also emerging for minimally invasive emphysema treatment [4]. Each time a patient undergoes a follow-up scan, the software could be used to generate a summary report for the physician. Feedback on the number and placement of devices, as well as treatment outcome, may prove beneficial to the physician for planning future studies; such information could also provide feedback to the device manufactures.

The automated techniques were designed to be generic, and thus hopefully extensible to other type of devices. The use of relevant, yet very simple control point features allows incorporation of a high degree of domain knowledge, while also allowing very efficient computation. Following the pre-processing, the template matching time takes approximately 10 seconds (3.2 GHz single processor, 2 GB RAM, and implemented in Java).

Improving the initial lung-hilum segmentation and pre-processing should improve the device detection performance. Inclusion of the entire hilum was done so that devices in collapse lung would not be missed. Since the vast majority of false positives were due to calcifications in the heart and vessels, the ability to segment only the collapse lung from heart would significantly reduce the number of false positives.

Future work can be done to apply device specific analysis on top of the detection process. Potential future applications include monitoring the placement of devices on multiple exams, labeling the implanted airways, and/or determining if all airways in the target treatment lobe were completely occluded.

Limitations of this research include testing on only one device; collecting more data for testing the technique on other devices is in process. An additional limitation is the number of samples

required for training; future work will focus on how similar, if not better, performance can be obtained with fewer training samples.

5. CONCLUSION

A method for detecting implanted endobronchial valves was developed and showed strong performance when tested. The detection method relied on a cascade classifiers and simple features computed at control points relative to a template of the device. These generic device modeling and detection techniques may be applicable to other devices and may be useful for monitoring the placement and function of implanted devices.

Acknowledgements

Funding for this work was provided by UC Discovery Grant IT106-10158.

Grant Support provided by University of California Discovery Grant IT106-10158 in collaboration with Broncus Technologies, Inc. and Emphasys Medical Inc.

References

1. Toma T, Hopkinson N, Hillier J, Hansell D, Morgan C, Goldstraw P, Polkey M, Geddes D. Bronchoscopic volume reduction with valve implants in patients with severe emphysema. *Lancet* 2003;361:931–933. [PubMed: 12648974]
2. Wan IYP, Toma TP, Geddes DM, et al. *Chest* 2006;129:518–526. [PubMed: 16537847]
3. Yim A, Hwong T, Lee T, Li W, Lam S, Yeung T, Hui D, Ko F, Sihoe A, Thung K, Arifi A. Early results of endoscopic lung volume reduction for emphysema. *J Thorac Cardiovasc Surg* 2004;127:1564–1573. [PubMed: 15173708]
4. Maxfield RA. New and Emerging Minimally Invasive Techniques for Lung Volume Reduction. *Chest* 2004;125:777–783. [PubMed: 14769765]
5. Harris E, McNair H, Evans P. Feasibility of fully automated detection of fiducial markers implanted into the prostate using electronic portal imaging: A comparison of methods. *Int J Radiation Oncology Biol Phys* 2006;66:1263–1270.
6. van der Weide R, Bakker C, Viergever M. Localization of Intravascular Devices With Paramagnetic Markers in MR Images *IEEE Trans. Med Imaging* 2001;20:1061–1071.
7. Brown MS, Goldin JG, McNitt-Gray MF, Greaser LE, Sapra A, Li KT, Sayre JW, Martin K, Aberle DR. Knowledge-based segmentation of thoracic CT images for assessment of split lung function. *Medical Physics* 2000;27:592–598. [PubMed: 10757610]
8. Abramson Y, Freund Y. SEmi-automatic VISuaL LEarning (SEVILLE): a tutorial on active learning for visual object recognition. *IEEE CVPR*. 2005
9. Viola P, Jones M. Rapid object detection using a boosted cascade of simple features. *IEEE CVPR*. 2001
10. Duda, RO.; Hart, PE.; Stork, DG. *Pattern Classification*. Vol. 2. John Wiley and Sons; New York, NY: 2001.
11. Ochs R, Goldin J, Abtin F, Kim HJ, Brown K, Batra P, Roback D, McNitt-Gray M, Brown M. Automated classification of lung bronchovascular anatomy in CT using AdaBoost. *Medical Image Analysis* 2007;11:315–324. [PubMed: 17482500]
12. Food and Drug Administration. Critical Path Opportunities List. 2006. http://www.fda.gov/oc/initiatives/criticalpath/reports/opp_list.pdf
13. Das KM, El-Menyar AA, Salam AM, Singh R, Dabdoob WAK, Albinali HA, Suwaidi JA. Contrast-enhanced 64-Section Coronary Multidetector CT Angiography versus Conventional Coronary Angiography for Stent Assessment. *Radiology* 2007;245:424–432. [PubMed: 17890354]

Figure 1.
CT image of a valve that migrated into the posterior parenchyma (top, bone window/level).
An image of another valve placed too distal for complete occlusion of all segmental airways
of a lobe in also shown (bottom, lung window/level).

Figure 2.

Illustration of the device model, control points, and control groups. A 2D device model and examples are shown for clarity, but a 3D device model with 58 control points rotated at 98 orientations is applied on the volumetric CT data.

Figure 3. Illustration of the underlying control points of the weak classifiers selected by the algorithm for the first level of the cascade classifier. Three of the points were on the outer ring of the device model while one was inside the device lumen.

Confidence Score Algorithm**Input:**

- Point template from AdaBoost cascade

Initialize confidence score, σ , to zero

For Each Control Point, $\bar{p}_0 \dots \bar{p}_n$

If $\bar{p}_i$ was model *on* the base or outer ring of the valve

If $I_{\max}(\bar{p}_i) > 700 HU$ **then** $\sigma = \sigma + 4$

Else If $I_{\max}(\bar{p}_i) > 400 HU$ **then** $\sigma = \sigma + 2$

Else If $I_{\max}(\bar{p}_i) > 200 HU$ **then** $\sigma = \sigma + 1$

Else ($\bar{p}_i$ was model *in the lumen or outside the valve*)

If $I_{\max}(\bar{p}_i) < -100 HU$ **then** $\sigma = \sigma + 1$

Return σ

Figure 4.

Algorithm for determining the confidence score for a device, which is based on the maximum attenuation feature at each control point and whether or not the control point was model as being part of a control group on the edge of the device or inside the device lumen. The confidence score is formed by the summation of all the scores for all control points.

Device Detection Implementation Algorithm**Input:**

- Test points remaining after pre-processing
- Templates at 98 unique orientations: $K_0 \dots K_{98}$
- Device specific AdaBoost cascade classifier

Initialize linked list to reference classified devices and their confidence score

For Each Test Point**For Each Template**

1. Translate template to test point
2. If point template is classified as device by AdaBoost cascade
 - a. Compute confidence score
 - b. Add point template (i.e., device) and confidence score to linked list

Finally

Accept devices with highest confidence score in descending order
Reject a device if it overlaps a previously accepted device

Figure 5. Implementation algorithm for processing through the test points and generating the final set of devices.

Figure 6.

Illustration of the detection process. The four levels of the cascade are used to reject non-device point templates; the majority of instances are rejected after the first cascade. More difficult instances are rejected in later cascades. The final set of devices is formed by accepting devices from the list in descending order according to their confidence score; overlaps are rejected.

Figure 7.

FROC curves for the device detection performance. Curve (a) is when the system did not know how many devices were implanted; curve (b) is when the system was told how many devices should be detected.

Figure 8. Example of two correctly detected devices in collapsed lung. Note the difficulty in differentiating collapsed lung from nearby tissue. The slice thickness/interval for this series was 0.6/0.5 mm.

Figure 9. Example of two missed device anterior to two found devices. The slice thickness/interval (1.25/0.7 mm) and orientation of the devices may have resulted in partial volume averaging, resulting in the devices being missed by the initial threshold.

Figure 10.

False positives were from bone, calcium deposits, and metal artifacts. The pattern of the streak artifact likely resulted in the false positive. The slice thickness/interval for this series was 1.5/0.8 mm.

Table 1

The classification conditions and training error of the four weak classifiers selected by the AdaBoost algorithm for the first level of the cascade classifier. Together, these four weak classifiers could reduce the initial number of negative samples by 90% during training. The point number corresponds to the illustration in Figure 3.

Point	Weak Classifier Condition	Training Error
1	Max > 320 HU	0.080
2	Max > 475 HU	0.092
3	Max > 460 HU	0.190
4	Max $\leq$ 150 HU	0.167

Table 2

Results of the training process; each additional level to the cascade caused an approximate 10-fold reduction in the number of negative samples. The number of weak classifiers used at each level as well as the specificity of the cascade classifier up to and including that level is also shown.

Cascade Level	Weak Classifiers (one per feature)	Negatives	Specificity
Initial	-	11433366	-
1	4	1109209	90.03%
2	8	114273	99.00%
3	16	33504	99.71%
4	32	1197	99.99%