Incremental Learning of Control Knowledge for Lung Boundary Extraction

Avishkar Misra, Arcot Sowmya and Paul Compton

School of Computer Science and Engineering
University of New South Wales
NSW, Australia
amisra@cse.unsw.edu.au
sowmya@cse.unsw.edu.au
compton@cse.unsw.edu.au

Abstract. The goal of this work was to develop an adaptable computer vision system that refines itself to the specific task of extracting lung boundary in High Resolution Computed Tomography (HRCT) scans. We have developed an incremental learning framework called ProcessRDR that allows the underlying procedures of a computer vision system to learn knowledge pertaining to their control. This approach to learning control knowledge provides a systematic mechanism to customisation of the procedures for a domain, whilst the system is in operation.

1 Introduction

Computer vision is generally formulated as a two step process. Firstly, image analysis for *feature extraction* processes the input image(s) to extract the feature of interest to the system. This is followed by *recognition and classification* of the extracted features according to a semantic model of what is expected in the scene or image(s).

The success of computer vision systems depends on the strength of feature extraction and classification processes. It is generally accepted that these processes work in tandem and complement each other. A complex and sophisticated feature extraction process would focus on the important features of interest and provide features that can accurately delineate the object in the scene. Therefore a simple classification process would do the job quite well. A system with a simplistic feature extraction, however, would require a significantly more complex and powerful classification system.

Computer vision has seen various approaches to improve both the classification and feature extraction stages. The initial focus was on improving the underlying algorithms or procedures used for feature extraction and using machine learning and pattern recognition to improve the classification processes. The learning for classification processes was directed to accommodate *domain knowledge* of what is expected in the scene.

However, by the last decade, computer vision experts had amassed a variety of underlying procedures and were faced with a new problem of how and where to use these procedures. Some procedures are only applicable under certain circumstances or for certain types of images. For example, image enhancement and restoration will work to improve a poor quality image but will adversely affect a high quality image in terms of loss of information. Other procedures, designed to be image independent, have to be specialised for a particular domain via their parameters.

The selection of optimal parameters for a procedure within a specific domain requires expertise in the field of computer vision and often in the specific domain of application. These experts draw upon their knowledge about computer vision and the domain of application to try and find the optimal parameters for the task. This process often ends up being rather an ad-hoc process of trial and error, at the end of which there is no guarantee that the chosen set of parameter values would be optimal in all situations.

This problem of learning *control knowledge* [20] to use vision procedures is complicated through the high degree of variability in the optimal solution for different applications. Not only do we wish to determine the optimal set of parameters for a procedure in a domain, but also how to combine a number of procedures to solve a larger and more complex feature extraction task.

In this paper we present a scheme for incremental learning of control knowledge for computer vision systems. We have built a system, which learns from the expert how to extract the lung boundary from High Resolution Computed Tomography images. The following section will provide an overview of relevant computer vision problems and the need for appropriate learning mechanisms. In section 3 we will introduce ProcessRDR and discuss a prototype lung boundary extraction system in section 4. The paper concludes in section 5.

2 Background

The advent of medical imaging technologies such as X-Ray, Magnetic Resonance Imaging, and Computed Tomography [41] provides doctors with a non-invasive alternative to looking inside a patient.

As a result Computer Aided Diagnosis (CAD) systems which use medical imaging technologies, have been an area of active research in the last two decades [1] [2] [4] [5] [6] [7]. [3] provides a good overview of the various techniques and developments within the field of Pulmonary (Lung) Imaging and Analysis, used by various CAD systems designed for the Lung.

The identification of lung boundary (pleura) within Lung CAD systems is an important step towards detecting diseases and abnormalities for a patient. An accurate delineation of the lung boundary is especially important when detecting diseases affecting the walls of the lung and its neighbouring regions.

A number of systems similar to [8] were developed to extract lung boundary, relying on computer vision procedures and hand-crafted heuristics that combine these procedures together. Noisy data or variations within the norm, can affect the results and true location of the lung boundary. Approaches like [11] and [12] look at ways to improve the feature extraction algorithms themselves. Here the responsibility of developing accurate feature extraction algorithms resides with the computer vision expert. In both cases the researchers have sought to improve the underlying feature extraction algorithm manually.

As the domain becomes more complex, hand-crafted feature extraction algorithms must be supplemented with heuristics and generalisations. Often these heuristics are developed on a trial-and-error basis and make amendments difficult. The development of these algorithms themselves is a difficult and time consuming process, with no assurance that the developed algorithm using an expert's control knowledge would be sufficiently flexible in all circumstances.

In order to provide the needed domain knowledge, [9] and [10] have used an anatomical model for segmentation. Their process involves generating a model of expectation, which is used to support their lung boundary segmentation algorithms. Though this approach incorporates greater degree of domain knowledge, the models developed for such systems need to address the large variability in the lung shape and size across patients and within a patient. For example the lung boundary in the top third of the lung has a remarkably different shape and size to the lung boundary in the bottom third of the lung. Even if the variability was accounted for during the stage of generating the model, we have to question the ability of these model-based approaches to accommodate and learn from their failures.

Such model-based computer vision systems have been of interest in medical and non- medical applications of computer vision. [13] [14] [15] [16] [17] [18]. The limitation of model-based techniques is that they are generally applied either at the classification stage or the later stages of feature extraction. Here these techniques do not learn control knowledge to guide the feature extraction and instead use domain knowledge to select features that best suit the model.

A number of researchers have looked at learning control knowledge, through Parameter Tuning and Task or Goal based expert shells for computer vision.

Work by [21], [22], [23], [24] and [25] learn the optimal set of parameters or values for the procedures, by using statistical measures of success. These approaches are not successful when we have limited amount of data, or when the measure of success cannot be defined in statistical terms, which is often the case.

In other approaches, [26], [27], [28], [29] and [30] have developed computer vision systems to try to learn from human experts and expect their teachers to manually describe the solution to a problem in a structured way. In [26] for example, the system requires a computer vision expert to explicitly define all concepts of a problem and its associated solution. This form of learning requires the Computer Vision expert to pre-empt all possible outcomes and define a complete solution to the problem. This is not only impractical for many domains but also fails to capture all possible knowledge for a particular domain.

Research in Knowledge Acquisition, [31] [32], has shown that though human experts maintain an internal structure to their knowledge, they lack the ability to communicate the complete knowledge in a structured way. Instead they can justify a decision made, by using their knowledge. So irrespective of the level of expertise of a person, he or she cannot completely articulate that knowledge.

Ripple Down Rules (RDR), [33] [34], propose an approach to knowledge acquisition that addresses this problem. RDR and its variants, [35] [36] [37], are Knowledge Bases, composed of a tree or chain of rules. Each rule and its parent rules define the context or conditions which must be met in order for the consequent of the child rule to classify or act on the queried input. If RDR's rules make a mistake, the expert can teach the system by creating a child rule as an exception to the rule where it failed. Here, the problem of experts articulating their knowledge is addressed, by asking the expert to justify his/her reasoning. This justification, which is intuitive and highly effective, forms the context of an RDR rule.

This approach can also extract tactical and non-factual knowledge which is even harder to articulate using traditional methods. RDR's ability to learn incrementally on a case-by-case basis is an additional benefit to applications where the data is sparse, making it difficult to learn using statistical measures or traditional machine learning.

A number of alternatives have been proposed to the aforementioned Single Classification RDR techniques. Nested-RDR[35] for example, works to describe complex concepts within separate RDRs of their own. Here each RDR learns to describe a specific concept and the connected RDRs work together to solve a more complex task.

Multiple Classification RDR (MCRDR)[36], provides the capacity of the same RDR structure to maintain differential classification for the same case. It allows for the possible alternatives in classification for a given case. Park et al, [38] have used MCRDR to carry out lung boundary extraction from X-ray images. The knowledge-based method used works only at the classification and region selection level. The MCRDR is never used to guide the underlying processing.

Evaluation of various lung extraction techniques has posed a dilemma of its own. Due to the high degree of variability across patients, there is no correct lung boundary against which segmented boundaries can be evaluated. There have been a number of attempts to describe and measure the level of successful boundary extraction, [39] [40], but there is no technique which provides an automated statistical measure. This problem is exacerbated when even radiologists themselves cannot always agree upon where the lung boundary should be drawn. This makes statistical evaluation of a lung boundary difficult, and we have to rely upon radiologists deeming a boundary to be acceptable or not.

3 ProcessRDR

We have already alluded to the difficulty of learning for feature extraction procedures within computer vision, which requires a great deal of expertise and refinement. The manual customisation of feature extraction algorithms for a particular domain can often lead the development in an ad-hoc way. Though some might argue that a formal approach to software engineering can address

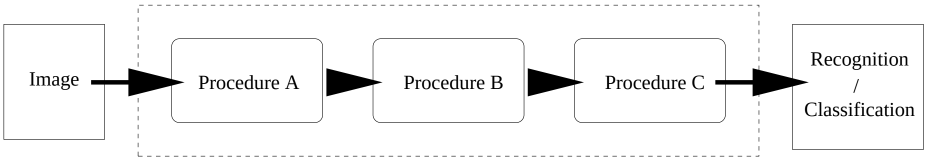

Fig. 1. Generalised Computer Vision system.

this ad-hoc nature, it can equally be argued that these do not work well unless the programmer has complete knowledge of all types of inputs within the domain. Considering that computer vision systems are developed by experts to solve problems in other highly specialised domains, it is unreasonable to expect the expert to understand all the nuances of that domain.

At the same time, work in Knowledge Acquisition [31] points out that even with complete knowledge, the programmer would have problems articulating a solution which cover all conceivable cases. Clearly, we need to develop computer vision systems that learn control knowledge to refine themselves during operation, and continue to improve beyond the initial training stage.

We propose a solution by extending RDR to feature extraction processes in computer vision. Here we are using RDR to learn control knowledge from the expert and subsequently use it to guide the underlying procedures during operations. Hence the so called ProcessRDR. RDR's learning mechanism provides a systematic approach to knowledge acquisition and maintenance, even when used in an ad-hoc method of operation.

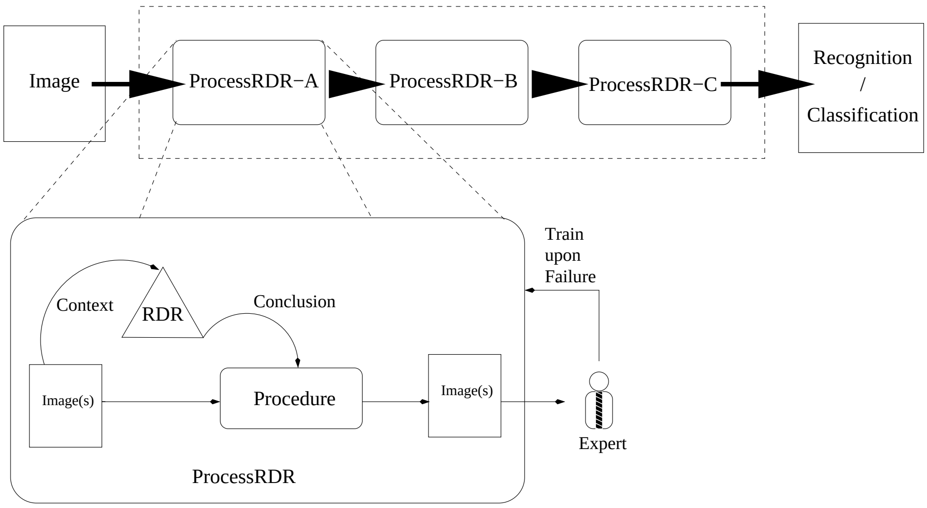

Fig. 2. ProcessRDR for a Computer Vision system.

A computer vision system, as shown in Figure 1, can be thought of as a sequence of procedures, connected up in a specific order by a programmer to solve the task at hand. ProcessRDR involves isolating important and configurable processes within a system, and attaching an RDR knowledge learner to it. Since complex computer vision systems have a number of procedures to carry out the necessary feature extraction, we would end up with a number of separate RDRs. A new image would pass through each of the ProcessRDR procedures in the sequence. Here the RDR controlled procedures carry out their specific processing on the image. For example in Figure 2, an image passing through ProcessRDR A would be transformed by the procedure. The transformed image would serve as the input for ProcessRDR B. The sequence of ProcessRDRs will conclude with the recognition and classification, which was the system's overall objectives.

RDR has been applied to a range of problems [37]. The closest of these is the use in configuration, also knowns as parameter optimisation. The difference with ProcessRDR is that, rather than producing some appropriate combination of parameter values in a static configuration, parameter values are selected for a sequence of processes and the output is the result of the processes, not the set of parameter values.

The learning of control knowledge involves learning of the optimal parameters for a specific case or domain, as well as the learning of sequencing of underlying computer vision procedures to solve a complex task. The ProcessRDR mechanism allows an expert to teach each of the procedures, the optimal operational control of that procedure or module which comprises of smaller atomic procedures.

The fine-grain control of learning with a specific task means that we can potentially teach a ProcessRDR using separate experts. This allows computer visions systems spanning a number of different domains to learn from each of the domain experts separately. For example, in a lung boundary extraction system the ProcessRDR would allow both computer vision experts and radiologists to collaborate on their domain-specific knowledge.

As the ProcessRDRs for a complex system are connected in a sequence, the conclusion for a case in a ProcessRDR, would effect the decisions and localised context of subsequent ProcessRDR. This is because the conclusion from one ProcessRDR serves as the input for the next ProcessRDR.

Therefore any correction made in ProcessRDR A in Figure 2, would warrant a re-evaluation of the cornerstone cases in ProcessRDR A as well as all the cornerstone cases in B and C. While any correction in ProcessRDR C will only require re-evalutation of local cornerstone cases.

4 ProcessRDR Application

In order to test our ideas, we applied ProcessRDR to address an important problem in medical imaging. The objective was to develop a system using ProcessRDR to learn control knowledge for the individual vision procedures to extract the lung boundary. We compared the performance of a non-ProcessRDR system against a system build using ProcessRDR. The non-ProcessRDR system was an existing simple lung boundary extraction system developed by Po, [8], known as Po Boundary Extraction. We compared the same cases and their results against an ProcessRDR Boundary Extraction system.

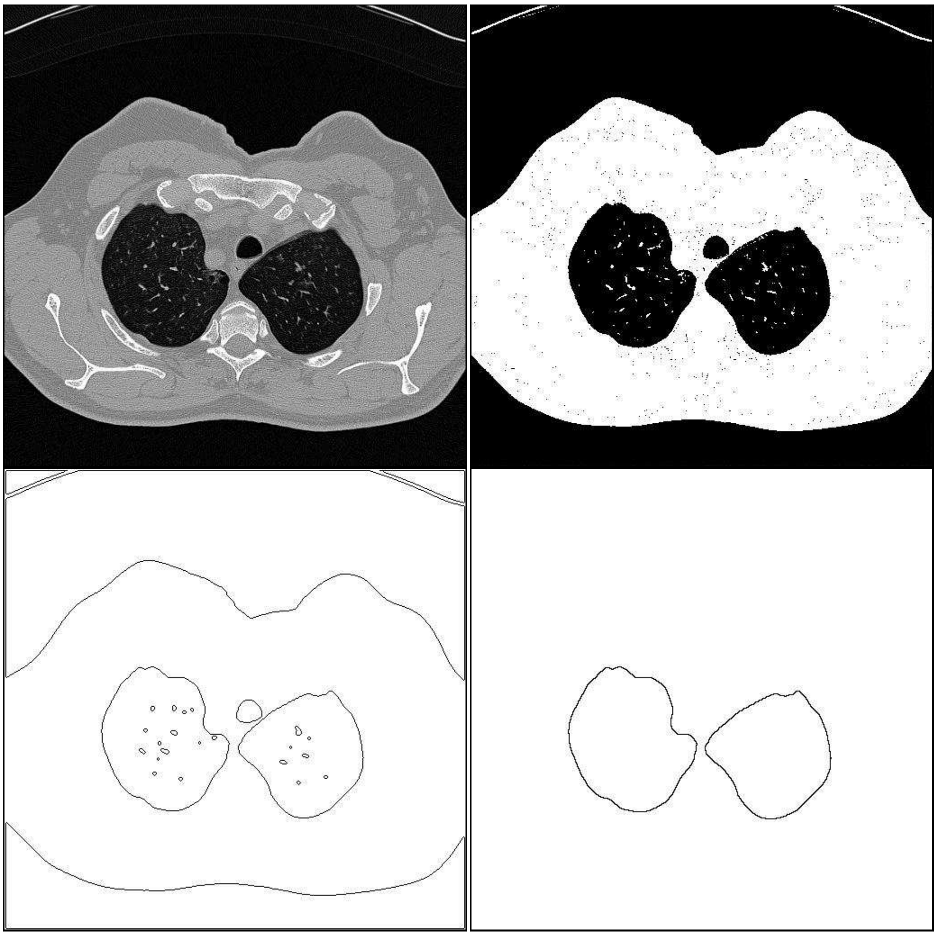

Fig. 3. A Generic Boundary Extraction Process (Top, Left to Right):(a) Original image, (b) after thresholding. (Bottom, Left to Right):(c) After morphological cleanup and (d) final lung boundary. Note that this sequence represents an example of expected behaviour, which Po Boundary Extraction does not achieve, and ProcessRDR Boundary Extraction converges towards.

4.1 Po Boundary Extraction

Po Boundary Extraction, hence forth known as Po-BE, uses standard Computer Vision procedures in an approach similar to existing boundary extraction systems mentioned earlier. The underlying algorithm has 3 main Vision Procedures and the results to each step can be seen in Figure 3.

1. Thresholding - carries out a grouping of similar pixel intensities and binary separation according to a defined threshold value. The pixels inside the lung have a lower pixel value compared to regions outside the lung. The thresholding process attempts to get the best separation between regions inside the lung and outside the lung as shown in Figure 3(b). The selection of the optimal threshold value can vary from image to image.
2. Morphological Operations - a series of erosion, dilation, opening or closing operators are applied to remove as much of the noise as possible, while preserving the integrity of the lung boundary. The order and number of times these operators are applied are important and can also vary from case to case. Once finished, an outline of the candidate regions is carried out as shown in Figure 3(c).
3. Connected-Component and Boundary Selection - Here the procedures find the most likely candidate for regions defining the boundary and present the extracted boundary. This is often very hard to define in explicit terms or even non-conflicting heuristics. Figure 3(d) shows an extracted lung boundary from the original input as shown in Figure 3(a).

Po developed this algorithm manually and through trial-and-error. In doing so, he produced optimal set of parameters for thresholding, sequences for morphological operators and boundary selection.

4.2 ProcessRDR Boundary Extraction

The ProcessRDR Boundary Extraction, hence forth known as ProcessRDR-BE, was developed to have the similar underlying procedures as Po-BE. The only difference was that we attached a RDR learner to each of the configurable procedures of the system. The procedure with the associated RDR forms a ProcessRDR bundle. The complete system is a connected sequence of ThresholderRDR, MorphologyRDR and RegionSelectionRDR.

The context for each of the ProcessRDRs to evaluate a case can be broken into the following categories:

1. Image properties - defines the properties of either the original Image or the image which is the immediate input to the ProcessRDR. It includes:
   - Statistical measures such as pixel intensity, mean and variance.
   - Objects of interest and their properties (i.e. number of objects)
   - Textures and region properties
2. HRCT or scan properties - defines the properties of the scan itself which are generally available from the HRCT header:

   - Patient orientation during scan. i.e. prone vs. supine
   - Slice location and range. i.e. image number 3 of 20.
   - Image reconstruction properties (scaling and algorithms)
3. Patient properties - defines the properties of the patient. This information is also available from the HRCT header, but may eventually include intermediate diagnosis as a result of querying external modules. Examples are:
   - Personal details such as age and sex.
   - Previously diagnosed disease or presence of other disease processes.

Though many of the context evaluations are common and can be treated in global fashion, we opted for a specific, localised evaluation within each RDR. The reason is that different ProcessRDRs might interpret the concepts in different terms. For example, the concept of mean in ThresholderRDR is seen as the mean of the pixel intensities, whilst the mean in RegionSelectionRDR refers to the mean area of candidate regions. Here the context evaluation changes due to the nature of the image data.

The conclusions within each ProcessRDR are the parameters to use for the processing and depend purely on the type of procedure we are dealing with. For example, ThresholderRDR's conclusion is the minimum and maximum thresholding values. The conclusions in MorphologyRDR, however are sequences and iterations of morphological operators.

The experts or users of the systems are responsible for training the RDR. Clearly there are two domains which interact and overlap - computer vision and medical imaging. In order to facilitate learning directly from users in medical imaging (i.e. radiologists) we developed appropriate graphical user interfaces, which allow the non- vision expert to articulate expected behaviour of the underlying procedure through a mouse. The designed interface allows a person who is not an expert in computer vision to actually define the context and processing conclusions for a procedure. For example, Figure 4 shows the interface to select the optimal threshold value. The user moves the scrollbar until the desired thresholding is achieved. We are able to build such interfaces for procedures that allow some form of mapping from the low- level procedure control/response to high-level visualisation.

4.3 Comparision/Results.

The ProcessRDR-BE system is still a sequence of vision procedures as with Po-BE and other similar techniques. The difference in ProcessRDR-BE is that each of these procedures acquire control knowledge of optimal operation within their domain, over the life-time of the application. ProcessRDR-BE will continue to expand its control knowledge and refine the underlying algorithm, directly from end user interactions.

Since there is no gold-standard in lung boundary evaluation, we could only evaluate the results of Po-BE and ProcessRDR-BE visually. Here the boundary is deemed acceptable or unacceptable, which automatically warrants a correction by the expert.

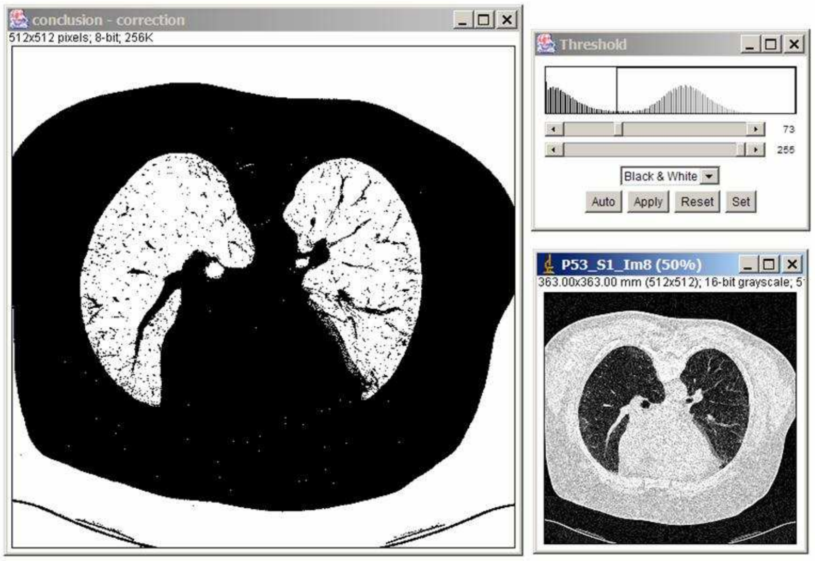

Fig. 4. GUI for teaching correct threshold value. Bottom-Right pane shows the original image. Left pane shows the affects of thresholding, which is controlled by the slider in the Top-Right.

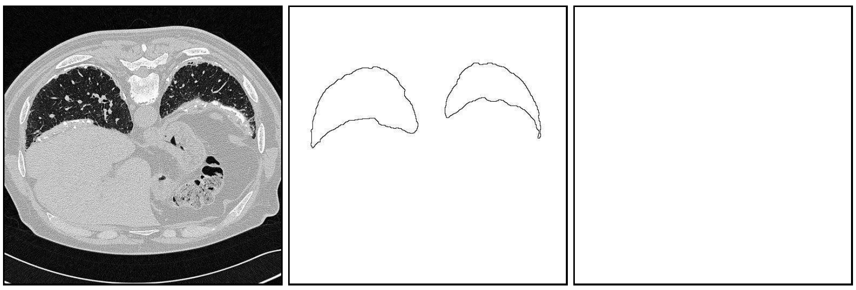

Fig. 5. Prior To Training (Left to Right): Original image, Po Boundary Extraction result and ProcessRDR Boundary Extraction with no result.

Prior to training, the ProcessRDR-BE's default rule settings were clearly not optimal and ProcessRDR-BE could produce no result, as shown in Figure 5. After training the performance of the ProcessRDR-BE started to converge and was comparable to Po-BE. Figure 6 shows a failure of both systems. Po-BE

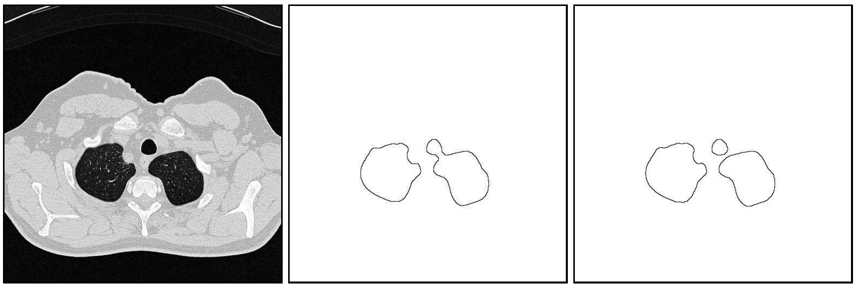

Fig. 6. Failure During Training (Left to Right): Original image, Po Boundary Extraction result and ProcessRDR Boundary Extraction result.

included a smaller circular region, known as the trachea, into the lung boundary. The trachea is often mistaken to be a part of the lung by Po-BE if it lies close to the boundary. This is due to selection of morphological operators and region selection procedures in the system. The ProcessRDR-BE's MorphologyRDR was able to keep the trachea isolated from the lung boundary, but the RegionSelectionRDR failed to eliminate the smaller region. The expert was able to define the appropriate rule to correct this behaviour as shown in Figure 7. Though it might seem that RegionSelection procedures within Po-BE could be directly improved, the reason why Po-BE failed to exclude the trachea from the boundary was a failure during Morphological cleanup of the image. Po-BE makes a single generalisation in order to accomodate cases where 'dilation' must be applied multiple times. ProcessRDR-BE gets around this problem by allowing the system to treat those cases differently.

5 Conclusion

We have presented an incremental learning technique for Computer Vision systems that uses Knowledge Acquisition to refine the control knowledge. In this work we have defined and used ProcessRDR, that is taught directly by the radiologists on how to extract the lung boundary.

The ProcessRDR framework allows for continual learning and refinement of control knowledge which guides the underlying procedures in a complex computer vision systems. Traditionally, these procedures are customised by experts via modification to the algorithm itself, development of supporting heuristics or

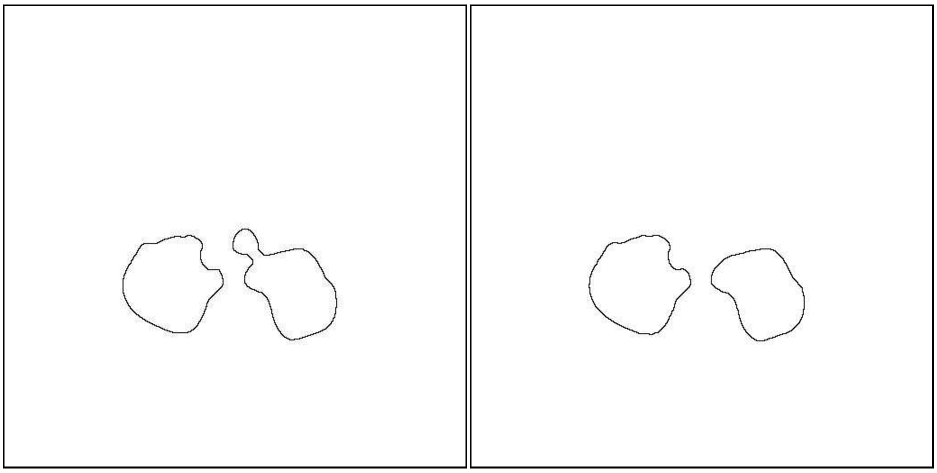

Fig. 7. After Correction of case shown in Figure 6 (Left to Right): Po-BE remains unchanged. ProcessRDR-BE's process is corrected.

a selection of parameters. The approach presented eliminates the problems associated with what is inherently an ad-hoc refinement process, and instead offers a structured approach to learning of knowledge to guide the processing.

ProcessRDR uses RDR's knowledge acquisition technique to learn control knowledge for computer vision procedures. The same technique can be applied to other forms of processing systems, that often require expert customization for a specific application.

We are currently applying ProcessRDR to more complex computer vision procedures to validate the flexibility of the ProcessRDR framework. In addition to this, we are also trying to address the issues of dependence between individual ProcessRDR and the significant number of cornerstone evaluations under some worst-case scenarios.

References

1. Ginneken, B.-v, Romeny, B.M., Viergever, M.A.: Computer-Aided Diagnosis in Chest Radiography:A Survey. IEEE Trans. Med. Imag, Vol. 20, No. 12, 2001
2. Shyu, C.-R., Brodley, C.E., Kak, A.C., Kosaka, A., Aisen, A.M., Broderick, L.S.: ASSERT: A Physician-in-the-Loop Content-Based Retrieval System for HRCT Image Databases., Computer Vision and Image Undersding: CVIU, (1999) Vol. 75. No. 1-2, p.111-132
3. Reinhardt, J.M., Uppaluri, R., Higgins, W.E., Hoffman, E.A.: Chapter 16: Pulmonary Imaging and Analysis Medical Image Processing and Analysis, Fitzpatrick, J.M, Sonka, M. p-1005, SPIE Press 2000.
4. Chabat, F., Hu, V.-P., Hansell, D.M., Yang, G.-Z.: ERS Transform for the Automated Detection of Bronchial Abnormalities on CT of the Lungs. IEEE Transactions on Medical Imaging, Vol 20. no 9, September 2001

5. Chiou, Y.S.P., Lure, Y.M.F., Ligomenides, P.A.: Neural network image analysis and classification in hybrid lung nodule detection (HLND) system. Proceedings of the IEEE-SP Workshop on Neural Networks for Signal Processing, 1993. p.517-526.
6. Linying, S., Sharp, B., Chibelushi, C.C.: Knowledge-Based Image Understanding, a rule-based production system for X-ray segmentation Proc. of the 4th International Conference on Enterprise Information Systems ICEIS 2002, 3-6 April 2003, 530-533 , 2002
7. Zhou, Z.-H., Jiang, Y.: Medical Diagnosis with C4.5 Rule Preceded by Artificial Neural Network Ensemble. IEEE Transactions on Information Technology in Biomedicine, 2003, 7(1): 37-42"
8. Chiu, P. T., Sowmya, A.: Lung Boundary Detection and Low Level feature extraction and analysis from HRCT images. VISIM: Information Retrieval and Exploration of Large Medical Image Collections, Utrecht, The Netherlands, Oct 2001 (held in conjunction with Fourth International Conference on Medical Image Computing and Computer-Assisted Intervention).
9. , Brown, M., Wilson, L., Doust, B., Gill, R., Sun, C.: Knowledge-Based Method for Segmentation and Analysis of Lung Boundaries in Chest X-Ray Images. Computerized Medical Imaging and Graphics, , vol. 22, pp. 463-477.
10. Zhang, L., Hoffman, E.A., Reinhardt, J.M.: Atlas-Driven Lung Lobe Segmentation in Volumetric X-ray CT Images. Proc. SPIE Conf. Medical Imaging (2003) San Diego, CA, pp309-319, Vol 5032.
11. Ukil, S., Reinhardt, J.M.: Smoothing Lung Segmentation Surfaces in 3D X-ray CT Images using Anatomic Guidance. Proceedings of SPIE Conference on Medical Imaging, 2004
12. Hu, H., Hoffman, E.A., Reinhardt, J.M.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. IEEE Trans. Medical Imaging, (2001), Vol. 20-6, pp-490-498, June 2001
13. Suetens, P., Fua, P., Hanson, A. J.: Computational strategies for object recognition. ACM Computing Surveys, (1992) 24(1):5–61
14. Arman, F., Aggarwal, J.K.: Model-based object recognition in dense-range images–a review. ACM Computing Surveys, (1993) Vol 26. No. 1, March 1993
15. Papageorgiou, C., Poggio, T.: A Trainable System for Object Detection. International Journal of COmputer Vision (2000) - 38(1), 15-33.
16. Kunz, D., Schilling, K., Ogtle, T.: A new approach for satellite image analysis by means of a semantic network. In W. F orstner and L. Plumer, editors, Semantic Modeling, pages 20–36, Basel, 1997. Birkh auser.
17. Kunz, D.: Investigation of synergy effects between satellite imagery and digital topographic databases by using integrated knowledge processing. International Archives of Photogrammetry and Remose Sensing, (1999) Vol. 32, Part 7-4-3 W6, pp. 144-151.
18. Bovenkamp, E.G.P., Dijkstra, J., Bosch, J.G., Reiber, J.H.: Multi-agent segmentation of IVUS images. Pattern Recognition, April 2004, Vol.37 Issue 4, pp.647-663.
19. Linying, S., Sharp, B., Chibelushi, C.C.: Knowledge-Based Image Understanding, a rule-based production system for X-ray segmentation. Proc. of the 4th International Conference on Enterprise Information Systems ICEIS 2002, 3-6 April 2003, 530-533.
20. Draper, B., Hanson, A., Riseman, E.: Knowledge-directed vision: control, learning and integration. Proc. of IEEE, vol. 84, n. 11, pp. 1625-1681, 1996"
21. Bhan, B., Lee, S.-K., Das, S.: Adaptive image segmentation using genetic and hybrid search methods. IEEE Transactions on Aerospace and Electronic Systems, (1994)

22. Bhanu, B., Peng, J.: Adaptive integrated image segmentation and object recognition. IEEE Trans. on Systems, Man and Cybernetics, Vol. 30, No. 4, pp. 427-441, November 2000.
23. Niemann, H., Fischer, V., Paulus, D., Fischer, J.: Knowledge Based Image Understanding by Iterative Optimization Advances in Artificial Intelligence, KI 96, Dresden, Görz, G., and Hölldobler, St., Vol. 1137, Springer, Berlin, pp287-301
24. Shekhar, C., Burlina, P., Moisan, S.: Design of SelfTuning IU Systems DARPA Image Understanding Workshop, volume 1, pp. 529536, New Orleans, LA, (May 1997).
25. Taniguchi, R., Amamiya, M., Kawaguichi, E.: Knowledge-based image processing system: IPSSENS-II. Proc. IEEE 3rd Int. Conf. On Image Processing and its Applications, Warwick, UK, 1989.
26. Clouard, R., Elmoataz, A., Porquet, C., Revenu, M.: Borg: A Knowledge-Based System for Automatic Generation of Image Processing Programs. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol 21, no. 2, pp. 128-144.
27. Clement, V., Thonnat, M.: Handling knowledge in image processing libraries to build automatic systems. International Workshop on Industrial Applications of Machine Intelligence and Vision, 1989.
28. Clement, V., Thonnat, M.: A Knowledge-Based Approach to Integration of Image Processing Procedures. ICVGIP - Image Understanding, volume 57 (1993), number 2. pp. 166-184
29. Lansky, A.L, Friedman, M., Getoor, L., Schmidler, S,. Short, N.Jr.: The Collage/Khoros Link: Planning for Image Processing Tasks. Integrated Planning Applications: Papers from the 1995, AAAI Spring Symposium. pp.67-76.
30. Wangehheim, A.V., Wagner, H., Comuello, W., Krechel, D. Richter, M.M., Conrad, P.: Cyclops - Expert System Shell for Development of Applications in Area of Medical Image Analysis Workshop 2000 of The German-Brazilian Cooperation Programme on Information Technology, 2000. , v.1. p.129-134.
31. Compton, P. J.,Jansen, R.: A philosophical basis for knowledge acquisition. 3rd European Knowledge Acquisition for Knowledge-Based Systems Workshop, Paris: 75-89 (1989).
32. Compton, P.: Insight and Knowledge AAAI Spring Symposium: Cognitive aspects of knowledge acquisition (1992) Boose, J., Clancey, W., Gaines, B., Rappaport, A., (Stanford University, 1992), pp 57-63.
33. Compton, P., Edwards, G., Kang, B., Lazarus, L., Malor, R., Menzies, T., Preston, P., Srinivasan, A., Sammut, C.: Ripple down rules: possibilities and limitations 6th Bannf AAAI Knowledge Acquisition for Knowledge Based Systems Workshop, Banff, (1991)
34. Preston, P., Edwards, G., Compton, P.: A 2000 Rule Expert System Without a Knowledge Engineer. Proceedings of the 8th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada. (1994).
35. Beydoun, G., Hoffmann, A.G.: NRDR for the Acquisition of Search Knowledge. Australian Joint Conference on Artificial Intelligence, pp177-186, 1997
36. Kang, B., Compton, P., Preston, P.: Multiple Classification Ripple Down Rules : Evaluation and Possibilities. Proceedings of the 9th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop, Banff, Canada,University of Calgary. (1995).
37. Compton, P., Richards, D.: Generalising Ripple-Down Rules. Knowledge Engineering and Knowledge Management: Methods, Models, Tools, Eds. R. Dieng; O. Corby, Juan-les-Pins France, 2-6 Oct. 2000, Springer, Berlin, pp380-386.

38. Park, M., Wilson, L., Jin, J.: Automatic Extraction of Lung Boundaries by a Knowledge-Based Method. Visualisation 2000, Eds. Sydney Australia, 4 Dec. 2000, Australian Computer Society, Sydney, pp11-16.
39. Chalana, V., Kim, Y.: A Methodology for Evaluation of Boundary Detection Algorithms on Medical Images. IEEE Transactions on Medical Imaging, Vol 16, No 5, October 1997, pp 642-652
40. Huang, Q., Dom, B.: Quantitative Methods of Evaluating Image Segmentation IEEE International Conference on Image Processing, 1995, pp 53-56
41. Kak, A.C., Slaney, M.: Principles of Computerized Tomographic Imaging, IEEE Press, 1988.