

An Automatically Generated Texture-based Atlas of the Lungs

Yashin Dicente Cid^{a,b}, Oula Puonti^{c,d}, Alexandra Platon^e, Koen Van Leemput^{c,f},
Henning Müller^{a,b}, and Pierre–Alexandre Poletti^e

^aUniversity of Applied Sciences Western Switzerland (HES–SO), Sierre, Switzerland

^bUniversity of Geneva, Switzerland

^cTechnical University of Denmark, Kgs. Lyngby, Denmark

^dCopenhagen University Hospital Hvidovre, Denmark

^eGeneva University Hospitals (HUG), Switzerland

^fMassachusetts General Hospital, Harvard Medical School, Charlestown, USA

ABSTRACT

Many pulmonary diseases can be characterized by visual abnormalities on lung CT scans. Some diseases manifest similar defects but require completely different treatments, as is the case for Pulmonary Hypertension (PH) and Pulmonary Embolism (PE): both present hypo- and hyper-perfused regions but with different distribution across the lung and require different treatment protocols. Finding these distributions by visual inspection is not trivial even for trained radiologists who currently use invasive catheterism to diagnose PH. A Computer-Aided Diagnosis (CAD) tool that could facilitate the non-invasive diagnosis of these diseases can benefit both the radiologists and the patients. Most of the visual differences in the parenchyma can be characterized using texture descriptors. Current CAD systems often use texture information but the texture is either computed in a patch-based fashion, or based on an anatomical division of the lung. The difficulty of precisely finding these divisions in abnormal lungs calls for new tools for obtaining new meaningful divisions of the lungs.

In this paper we present a method for unsupervised segmentation of lung CT scans into subregions that are similar in terms of texture and spatial proximity. To this extent, we combine a previously validated Riesz-wavelet texture descriptor with a well-known superpixel segmentation approach that we extend to 3D. We demonstrate the feasibility and accuracy of our approach on a simulated texture dataset, and show preliminary results for CT scans of the lung comparing subjects suffering either from PH or PE. The resulting texture-based atlas of individual lungs can potentially help physicians in diagnosis or be used for studying common texture distributions related to other diseases.

Keywords: Lung atlas, 3D texture, Riesz-wavelet, Supervoxels

1. INTRODUCTION

Most pulmonary diseases present visual abnormalities on the lung parenchyma in CT scans, such as hypo- and hyper-perfused regions or areas with unusual patterns. Pulmonary Hypertension (PH) and Pulmonary Embolism (PE) are two vascular pathologies presenting similar symptoms and visual defects on a CT scan but requiring very different treatments.^{1,2} Their similarities hinder their differentiation by health professionals,³ who currently rely on an invasive catheterism procedure for diagnosis.⁴ A PE usually presents one or more large hypo-perfused regions due to big emboli whereas PH is characterized by mosaic patterns with smaller hypo-perfused regions produced by micro-emboli. These distributions are not homogeneous across patients and their visual detection on a 3D volumetric scan is challenging for the human eye. However, hypo-perfused regions and other visual defects on the parenchyma can be characterized using computational texture features.⁵ Thin slices of tomographic scans have made it possible to represent many medical structures as isotropic solid 3D textures. Visualizing solid 3D texture is difficult and full 3D analysis is also not trivial even though many approaches exist.⁶ Most of the

Further author information: Send correspondence to Y.D.C. Email: yashin.dicente@hevs.ch

work using texture to detect affected patients has focused on quantifying the proportion of abnormal texture present.^{7,8} However, in the case of PE and PH, it is not the proportion that differentiates them but rather the spatial distribution across the lung of these textures.

An atlas is a localization system dividing the object of interest into regions with some common characteristics, like anatomical function, spatial position, or visual appearance.⁹ Many automatically generated atlases of the lungs have been proposed, mainly focusing on segmenting anatomical parts,¹⁰ but also creating geometrical regions.¹¹ The hypo-perfused regions in PE patients usually follow the lobe divisions. Using a lobe-based atlas for the analysis of these patients seems adequate but it may not work for PH cases where the defects are found in the entire parenchyma. Moreover, automatic lobe segmentation of the lungs is not an easy task on abnormal cases.¹² In,¹³ the authors showed that dividing the lung using a geometrical atlas and analyzing the relations of its regions helped to distinguish between PE and PH. However, these regions may contain more than one texture and the analysis of the entire region may hide important defects. On the other hand, Depeursinge et al.¹⁴ proposed an artificial division of the lungs based on classifying the 2D textures present in each CT slice among five predefined 2D textures (healthy, emphysema, ground glass, fibrosis and micronodules). While this approach helped the diagnosis of patients with interstitial lung diseases, none of these textures are characteristic of PE or PH, and the resulting atlas relied on previous knowledge (manual annotations) of the texture to detect.

In this work we present an automatically generated texture-based atlas of the lungs without relying on a priori knowledge of the textures to be found. As texture features we used the previously validated locally-oriented 3D Riesz-wavelet transform.¹⁵ Finally, we created the different regions of the atlas using a 3D extension of a superpixel¹⁶ algorithm (supervoxels) based on spatial and texture information. In addition to this atlas, we make a synthetic solid 3D texture database available to the scientific community for training the segmentation algorithms. Preliminary experiments on chest CT scans show the contribution of our trained algorithm to create a meaningful texture-based atlas of the lungs, capable of characterizing PH and PE patients.

2. METHODS

This section describes the details concerning the datasets and the techniques employed to carry out the experiments.

2.1 Databases

Two databases were used: a database of synthetic textures for training and validation purposes, and a database containing real CT scans of PE and PH patients. The synthetic database was created based on the RFAI database* (Reconnaissance de Formes, Analyse d'Images) of 3D synthetic textured images.¹⁷ The RFAI database also includes a set for texture segmentation, the *RFAI-segmentation* set. Our new database (available at[†]) contains 75 textured volumes of $128 \times 128 \times 128$ voxels and can be seen as an extension to the RFAI-segmentation set. Each volume is composed of a background texture and n insertions, where n varies from one to five and there are 15 volumes for each n -insertion class. Each insertion contains a different texture with random shape, size and location. The possible shapes are limited to sphere, cube, or nodule, allowing overlap between them and assuring a minimum size of 1,000 voxels. This database was divided into training and test sets using 50 and 25 volumes respectively. Both subsets contain an equal number of volumes of each n -insertion class, ten in the training set and five in the test set. Figure 1 shows one example of each n -insertions class.

The database of chest CT scans was provided by the radiology department of the University Hospitals of Geneva. It is composed of chest CT scans at 80 KeV of 18 PE patients, 45 patients with proven PH by a catheterism procedure, and 35 Control Cases (CC). The CC were patients suspected to have PE and diagnosed negatively after the visual inspection of the CT scan by an expert radiologist. The volumes were previously resampled to 1 mm isometric voxels and the lungs were automatically segmented using the algorithm presented in.¹⁸ Figure 2 shows one CT slice per group (CC, PE, and PH) with the corresponding lung segmentation.

*http://www.rfai.li.univ-tours.fr/PublicData/3D_Textures/3Dsynthetic_images_database.html, as of December 4, 2017.

[†]<http://publications.hevs.ch/index.php/publications/show/2331>, as of December 4, 2017.




Figure 1. 3D visualization of one example for each n -insertion class (from one to five insertions) of our synthetic database with the corresponding ground truth.




Figure 2. Coronal slices of three patients, one of each group (CC, PE, and PH), including in blue the automatically generated lung segmentation. The slices, acquired at 80 KeV, are shown in the range $[-1024, 300]$ HU, covering the full range of HU values inside the lungs. Using this standard range, the hypo- and hyper-perfused regions can only be identified in the PE case (central image).

2.2 Texture features

As texture descriptors we used the locally-oriented 3D Riesz-wavelet transform introduced by Dicente et al.¹⁵ The configuration chosen for this work obtained one of the best results for texture classification on the RFAI database: the 3rd order Riesz transform, with 4 scales and the local-orientation method based on the 1st order Riesz filters (see¹⁵). This configuration provides 40-dimensional feature vectors for each voxel of the volume.

Feature normalization: In order to handle the different range of energy of each wavelet scale, the ten components of each scale were divided by the mean energy in the corresponding scale:

$$f(v) = \left(\frac{|\mathcal{R}_1^1(v)|}{\mu_I(\mathcal{R}^1)}, \dots, \frac{|\mathcal{R}_{10}^1(v)|}{\mu_I(\mathcal{R}^1)}, \frac{|\mathcal{R}_1^2(v)|}{\mu_I(\mathcal{R}^2)}, \dots, \frac{|\mathcal{R}_{10}^4(v)|}{\mu_I(\mathcal{R}^4)} \right), \quad (1)$$

where $\mathcal{R}_i^s(v)$ is the response to the Riesz filter i at scale s in a voxel $v \in I$, and $\mu_I(\mathcal{R}^s)$ is the mean of the energies of all the filters at scale s over the set of voxels in I . In the case of the synthetic dataset, I corresponded to the set of voxels of each volume independently, while for the CT images, I referred to all the voxels of all patients, obtaining a global normalization over the full CT dataset.

2.3 Image segmentation using supervoxels

To segment a multi-channel Riesz feature image $\mathbf{X}_F = (\mathbf{x}^1, \dots, \mathbf{x}^F)$, where $\mathbf{x}^f = (x_1^f, \dots, x_I^f)^T$ denotes a single feature channel with I voxels, we rely on a supervoxel approach akin to¹⁹ where the voxels are clustered based on their similarity in feature space and spatial proximity on the image grid. Given the full feature image $\mathbf{X} = (\mathbf{X}_F, \mathbf{X}_S)$, where $\mathbf{X}_S = (\mathbf{x}^1, \mathbf{x}^2, \mathbf{x}^3)$ collects the spatial coordinates, the clustering is obtained by minimizing:

$$\mathcal{D} = \lambda \mathcal{D}_f + (1 - \lambda) \mathcal{D}_s, \quad (2)$$

where $\mathcal{D}_f = \sum_k \sum_{i \in C_k} \|\mathbf{x}_i^f - \boldsymbol{\mu}_k^f\|_2$, and $\mathcal{D}_s = \sum_k \sum_{i \in C_k} \|\mathbf{x}_i^s - \boldsymbol{\mu}_k^s\|_2$. Here, \mathcal{D}_f and \mathcal{D}_s denote the within-cluster feature space and image grid distances respectively, $\boldsymbol{\mu}_k^f$ and $\boldsymbol{\mu}_k^s$ are the centers of cluster C_k for the features and grid distances, and λ is the parameter controlling the relative contribution of the features against spatial proximity. The cluster means and assignments are easily found using the well-known K-means algorithm. As a final post-processing step after the segmentation, small supervoxels are fused with adjacent ones to ensure each supervoxel contains a minimum of 1,000 voxels.

3. EXPERIMENTAL RESULTS

This section first describes the quantitative results on the synthetic dataset, and next the qualitative results on real CT scans of patients and healthy controls.

3.1 Results on the synthetic dataset

The supervoxel algorithm is initialized by a user-defined step size δ that determines the initial position and size of the supervoxels. Furthermore, the user needs to define the feature and spatial weight λ . The K-means clustering is run until the largest spatial change, relative to the previous iteration, in the cluster centers falls below one-thousandth of the initial step size or until 100 iterations are completed. We find the optimal values of δ and λ by grid-search on the training dataset. The optimization is based on the state-of-the-art performance metrics Boundary Recall ($BRec$) and the Undersegmentation Error ($UErr$).²⁰ Figure 3 shows with an example how these metrics are affected when δ and λ vary, along with the resulting supervoxelization. The results obtained in the training set for the different values of δ and λ are shown in Figure 4. From now on, the performance of a set of parameters is defined as the average performance on the entire set of volumes. Choosing the best parameters for the supervoxelization algorithm is not trivial since the best performance for each metric was obtained with different values of δ and λ . We decided to choose the combination of parameters that minimizes the relative loss of performance (L^r) in both measures, i.e.:

$$\arg \min_{\delta, \lambda} (L_{BRec}^r(\delta, \lambda) + L_{UErr}^r(\delta, \lambda)), \quad (3)$$

where

$$L_{BRec}^r(\delta, \lambda) = \frac{\max_{\delta, \lambda} (BRec) - BRec(\delta, \lambda)}{\max_{\delta, \lambda} (BRec)}, \text{ and } L_{UErr}^r(\delta, \lambda) = \frac{UErr(\delta, \lambda) - \min_{\delta, \lambda} (UErr)}{\min_{\delta, \lambda} (UErr)}.$$

BRec and UErr refer in this case to the average BRec and UErr in the training set, respectively. Using this method the best combination was $\delta = 20$ and $\lambda = 0.8$. Table 1 shows the results for the training, test and RFAI-segmentation sets when using these optimal parameters.

Table 1. Quantitative results obtained in the training, test and RFAI-segmentation sets when using the best set of parameters ($\delta = 20$ and $\lambda = 0.8$).

Set	BRec (%)		UErr (%)	
	Mean \pm Std	Median [P25, P75]	Mean \pm Std	Median [P25, P75]
Training	81.35 \pm 3.03	82.99 [78.23, 86.30]	30.31 \pm 8.90	27.96 [21.75, 39.81]
Test	83.79 \pm 4.11	83.88 [79.34, 86.47]	31.08 \pm 10.23	30.34 [19.23, 42.05]
RFAI-segmentation	81.19 \pm 6.92	80.57 [77.79, 87.14]	12.19 \pm 7.01	10.04 [8.96, 14.11]

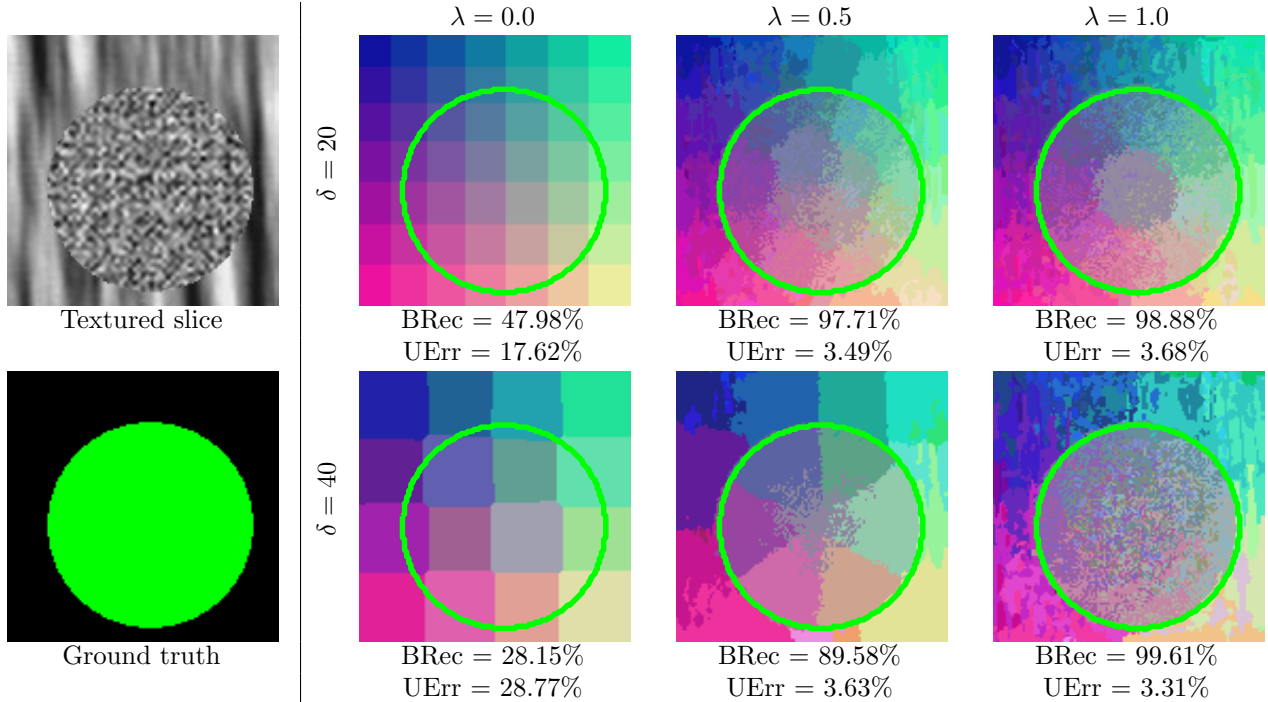


Figure 3. Example of supervoxelization when varying the initial step size (δ) and the feature-spatial weight (λ) (see Equation 2). The left-top and left-bottom images correspond to the textured 2D slice and the ground truth, respectively, of the 1-insertion class volume used in this example.

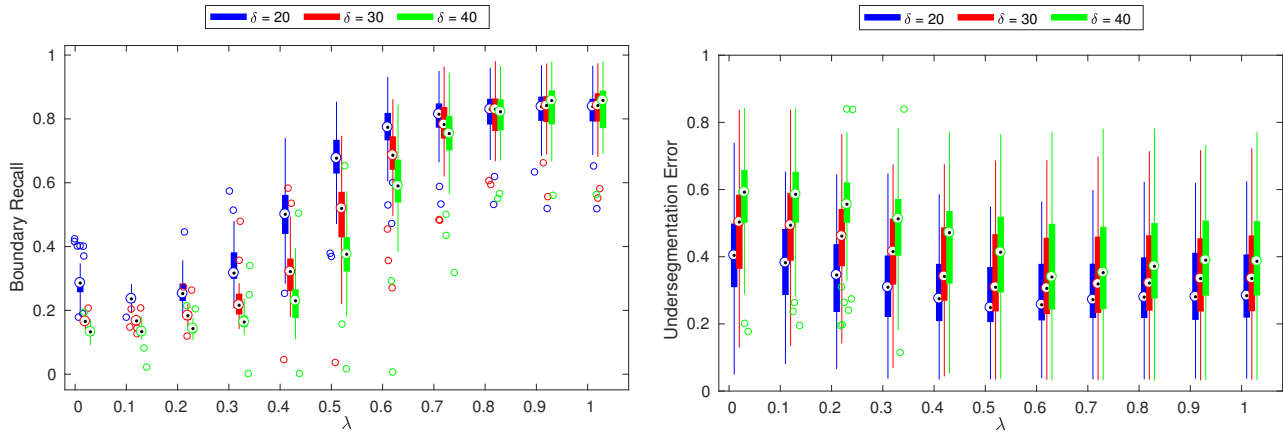


Figure 4. Performance measures obtained in the training set when δ and λ vary. The highest BRec was obtained when $\delta = 40$ and $\lambda = 1.0$, while the lowest UErr was at $\delta = 20$ and $\lambda = 0.5$. The optimal parameters considering both measures were found using Equation 3.

3.2 Results on the CT scans

Using the same optimal parameters found on the synthetic training set (see Section 3.1), the algorithm was run on the CT lung data to obtain a supervoxel-based atlas for each patient. Each supervoxel or region of this atlas can be seen as a texture unit with a specific feature signature corresponding to the average Riesz energy in the region. Let us call this signature $E_{v_k}(\mathcal{R}) \in \mathbb{R}^{40}$, being v_k the supervoxel k . In Figure 5 the average distribution of the norm is shown per scale of these signatures ($\|E_{v_k}(\mathcal{R}^s)\|_2$) over the different groups of patients (CC, PE, and PH). An initial analysis of these distributions suggest differences in the atlases of the patients from the 3 groups. Since the clearest difference between the classes is obtained at the first scale, we show the atlases using this scale in Figure 6. The three patients shown in the figure present scattered areas with low norm of

the Riesz energy in light blue. While in the CC these areas are small and almost not present, the corresponding region in the PE patient is clearly bigger and the PH patient shows medium size regions scattered in the entire parenchyma. Moreover, the areas corresponding to the main vessels are colorized differently in the three cases. Both PH and CC have high norm in these regions while the PE patient seems to have a lower norm.

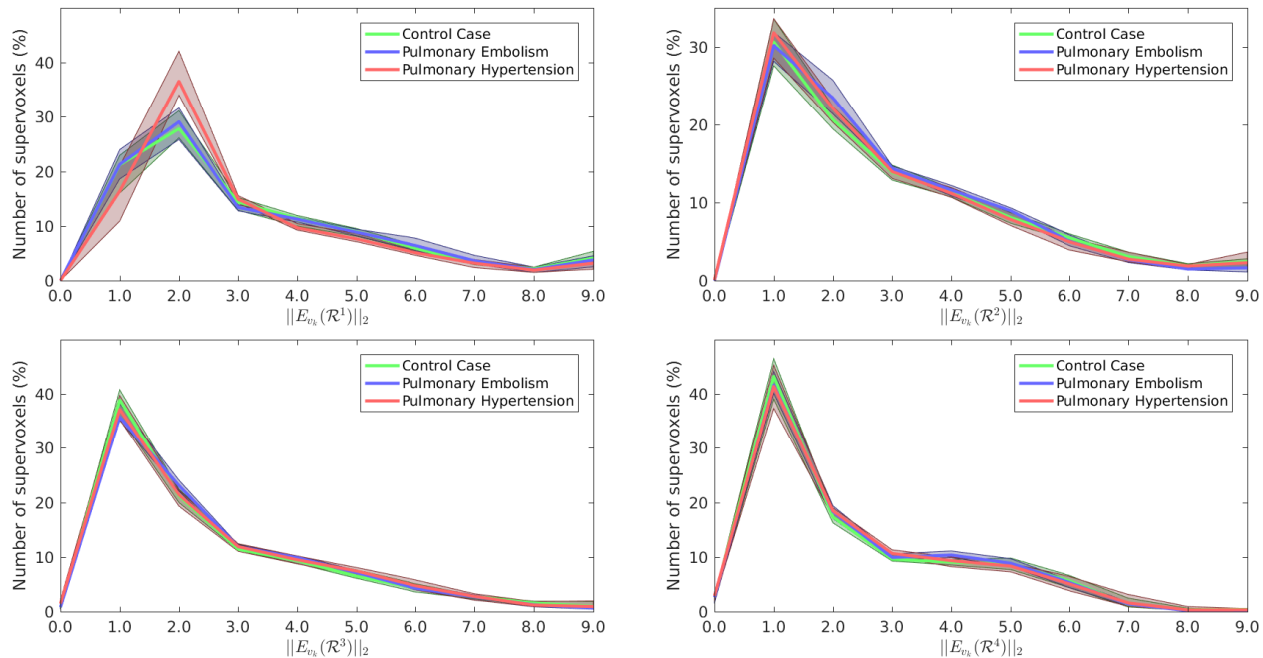


Figure 5. Distribution inside the atlases of the norms at different scales of the Riesz energies of each supervoxel. The bold line correspond to the median over all the patients of the corresponding group (CC, PE, and PH), while the shaded area corresponds to the region covered by the percentiles 33 and 66.




Figure 6. Resulting atlases of the CT slices shown in Figure 2. The color corresponds to the norm of the Riesz energy in each supervoxel ($\|E_{v_k}(\mathcal{R}^1)\|_2$).

4. DISCUSSION

To the best of our knowledge, the RFAI database is the largest database of 3D synthetically textured images for the time being. Among the different sets, it also includes a set of volumes for texture segmentation (RFAI-segmentation). However, this set only includes 15 volumes, not being sufficient for splitting it into training and test sets. A larger dataset is clearly required for our objectives. Moreover, the shape of the insertions in the

RFAI-segmentation set was limited to sphere and cube, and the size of these was significantly bigger than in the dataset created for this work. The results obtained on all three synthetic sets were similar, showing that the training set generated was suitable to train the supervoxelization algorithm. The UErr obtained on the RFAI dataset is significantly better. This improvement is based on the larger size of the insertions in the RFAI-segmentation set with respect to our sets. The minimum size of a relevant nodule in the lung parenchyma tends to be of 1 cm^3 . Aiming to emulate a real environment, the insertions of the presented dataset had a minimum size of 1,000 voxels, considering a standard voxel resolution of $1 \times 1 \times 1 \text{ mm}$, that matches the resolution of the resampled CT images used in this work.

Larger values of λ obtained better results, however, the supervoxelizations were more scattered, providing non-uniform regions more difficult to interpret by visual inspection (see Figure 3). For providing more meaningful atlases, the supervoxels should be grouped and/or cleaned by applying morphological operations. A measure providing information about the homogeneity and shape of the supervoxels could have helped in choosing a better set of parameters for the final atlas of the lungs.

Only PH patients were distinguishable with the atlas distribution shown in Figure 5. The many small changes in the parenchyma due to the micro-emboli of PH patients may explain these differences. CC and PE patients had similar distributions. However, it is worth to recall that the CC were patients suspected to have PE (i.e. not fully healthy patients) and may present similar distributions. More detailed analysis of the atlases can be done, since these preliminary differences were found by only analyzing the number of supervoxels at different scales and by using only the norm that may group regions with very different textures. In addition, these atlases can be used to build a *map of textures* for different diseases by averaging over patients, i.e., looking at the spatial distribution of the textures. Finally, since the normalization of the texture features was done based on all patients in the CT dataset, the resulting atlases depend directly on the set of patients used. All patients were visually inspected by expert radiologists and none of them contained any other major defect due to a non-related disease. The defects in the parenchyma present in the patients used in this work were limited to hypo- and hyper-perfused regions. Thus, the atlases built are specialized to detect these defects in the perfusion. Different sets of patients provide atlases highlighting other relevant regions.

5. CONCLUSIONS

In this work we presented a novel automatically generated atlas of textures of the lungs using a supervoxelization algorithm on local 3D texture descriptors. We created a database to optimize the parameters of the supervoxelization algorithm, providing a quantitative framework. The initial atlases already show promising results on characterizing PH versus PE patients and control cases, thus enabling its use as a lung descriptor to feed into machine learning techniques. Nonetheless, better refinement on the regions is required to provide a more manageable atlas in terms of shape and size of the regions. We believe such an atlas is helpful in CAD systems and, if included in diagnosis, may help in the description and detection of pulmonary diseases presenting abnormalities in the parenchyma. Moreover, the same concept of a texture-based atlas can also be applied to many other organs.

ACKNOWLEDGMENTS

This work was partly supported by the Swiss National Science Foundation in the PH4D project (grant agreement 320030-146804).

REFERENCES

- [1] Schwickert, H. C., Schweden, F., Schild, H. H., Piepenburg, R., Düber, C., Kauczor, H.-U., Renner, C., Iversen, S., and Thelen, M., “Pulmonary arteries and lung parenchyma in chronic pulmonary embolism: preoperative and postoperative CT findings.,” *Radiology* **191**(2), 351–357 (1994).
- [2] Tuder, R. M., Archer, S. L., Dorfmueller, P., Erzurum, S. C., Guignabert, C., Michelakis, E., Rabinovitch, M., Schermuly, R., Stenmark, K. R., and Morrell, N. W., “Relevant issues in the pathology and pathobiology of pulmonary hypertension.,” *Journal of the American College of Cardiology* **62**(25 SUPPL.), D4–D12 (2013).

- [3] Farber, H., “Pulmonary circulation: Diseases and their treatment. 3rd edition.,” *European Respiratory Review* **21**(123), 78–78 (2012).
- [4] Kim, N. H., Delcroix, M., Jenkins, D. P., Channick, R., Dartevelle, P., Jansa, P., Lang, I., Madani, M. M., Ogino, H., Pengo, V., and Mayer, E., “Chronic thromboembolic pulmonary hypertension,” *Journal of the American College of Cardiology* **62**(25 SUPPL.), D92–D99 (2013).
- [5] Depeursinge, A., Chin, A. C., Leung, A. N., Terrone, D., Bristow, M., Rosen, G., and Rubin, D. L., “Automated classification of usual interstitial pneumonia using regional volumetric texture analysis in high-resolution CT,” *Investigative Radiology* **50**, 261–267 (Apr. 2015).
- [6] Depeursinge, A., Al-Kadi, O. S., and Mitchell, J. R., [*Biomedical Texture Analysis: Fundamentals, Applications and Tools*], Elsevier-MICCAI Society Book series, Elsevier (Oct. 2017).
- [7] Burner, A., Donner, R., Mayerhoefer, M., Holzer, M., Kainberger, F., and Langs, G., “Texture bags: Anomaly retrieval in medical images based on local 3D–texture similarity,” in [*Medical Content-based Retrieval for Clinical Decision Support*], Greenspan, H., Müller, H., and Syeda-Mahmood, T., eds., *Lecture Notes in Computer Sciences* **7075**, 116–127, Springer, Heidelberg, Germany (September 2012).
- [8] Depeursinge, A., Foncubierta-Rodríguez, A., Vargas, A., Van De Ville, D., Platon, A., Poletti, P.-A., and Müller, H., “Rotation-covariant texture analysis of 4D dual-energy CT as an indicator of local pulmonary perfusion,” in [*IEEE 10th International Symposium on Biomedical Imaging, ISBI 2013*], 149–152, IEEE (Apr. 2013).
- [9] Li, B., Christensen, G. E., Hoffman, E. A., McLennan, G., and Reinhardt, J. M., “Establishing a normative atlas of the human lung: Computing the average transformation and atlas construction,” *Academic Radiology* **19**(11), 1368–1381 (2012).
- [10] Rikxoort, E. M. v. and Ginneken, B. v., “Automated segmentation of pulmonary structures in thoracic computed tomography scans: a review,” *Physics in medicine and biology* **58**(17), R187 (2013).
- [11] Depeursinge, A., Zrimec, T., Busayarat, S., and Müller, H., “3D lung image retrieval using localized features,” in [*Medical Imaging 2011: Computer-Aided Diagnosis*], **7963**, 79632E, SPIE (Feb. 2011).
- [12] Mansoor, A., Bagci, U., Foster, B., Xu, Z., Papadakis, G. Z., Folio, L. R., Udupa, J. K., and Mollura, D. J., “Segmentation and image analysis of abnormal lungs at ct: current approaches, challenges, and future trends,” *RadioGraphics* **35**(4), 1056–1076 (2015).
- [13] Dicente Cid, Y., Müller, H., Platon, A., Janssens, J.-P., Lador, F., Poletti, P.-A., and Depeursinge, A., “A lung graph-model for pulmonary hypertension and pulmonary embolism detection on DECT images,” in [*MICCAI Workshop on Medical Computer Vision: Algorithms for Big Data, MCV 2016*], (2016).
- [14] Depeursinge, A., Vargas, A., Gaillard, F., Platon, A., Geissbuhler, A., Poletti, P.-A., and Müller, H., “Content-based retrieval and analysis of HRCT images from patients with interstitial lung diseases: a comprehensive diagnostic aid framework,” in [*Computer Assisted Radiology and Surgery (CARS)*], (June 2010).
- [15] Dicente Cid, Y., Müller, H., Platon, A., Poletti, P.-A., and Depeursinge, A., “3-D solid texture classification using locally-oriented wavelet transforms,” *IEEE Transactions on Image Processing* **26**, 1899–1910 (Apr. 2017).
- [16] Ren, X. and Malik, J., “Learning a classification model for segmentation,” in [*Proceedings of the Ninth IEEE International Conference on Computer Vision - Volume 2*], *ICCV '03*, 10–, IEEE Computer Society, Washington, DC, USA (2003).
- [17] Paulhac, L., Makris, P., and Ramel, J.-Y., “A solid texture database for segmentation and classification experiments,” in [*In 4th International Conference on Computer Vision Theory and Applications*], 135–141 (Feb. 2009).
- [18] Dicente Cid, Y., Jimenez-del-Toro, O., Depeursinge, A., and Müller, H., “Efficient and fully automatic segmentation of the lungs in CT volumes,” in [*Proceedings of the VISCERAL Challenge at ISBI*], Orcun Goksel, Jimenez-del-Toro, O., Foncubierta-Rodríguez, A., and Müller, H., eds., *CEUR Workshop Proceedings*, 31–35 (Apr 2015).
- [19] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Süsstrunk, S., “SLIC superpixels compared to state-of-the-art superpixel methods,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**, 2274–2282 (Nov. 2012).

- [20] Neubert, P. and Protzel, P., “Superpixel benchmark and comparison,” in [*Proc. Forum Bildverarbeitung*], 1–12 (2012).