Published in final edited form as:
Acad Radiol. 2014 December ; 21(12): 1530–1541. doi:10.1016/j.acra.2014.07.017.

Extending Semi-Automatic Ventilation Defect Analysis for Hyperpolarized ^{129}Xe Ventilation MRI

Mu He, MS[1], **Suryanarayanan S. Kaushik, MS**[1,2], **Scott H. Robertson, MS**[1,4], **Matthew S. Freeman, MS**[1,4], **Rohan S. Virgincar, MS**[1,2], **H. Page McAdams, MD**[5], and **Bastiaan Driehuys, PhD**[1,2,3,4]

[1]Center for In Vivo Microscopy, Duke University Medical Center, Durham, NC, USA

[2]Department of Biomedical Engineering, Duke University, Durham, NC, USA

[3]Department of Radiology, Duke University Medical Center, Durham, NC, USA

[4]Medical Physics Graduate Program, Duke University, Durham, NC, USA

[5]Department of Cardiothoracic Imaging, Duke University Medical Center, Durham, NC, USA

Abstract

Rationale and Objectives—Clinical deployment of hyperpolarized (HP) ^{129}Xe MRI requires accurate quantification and visualization of the ventilation defect percentage (VDP). Here, we improve the robustness of our previous semi-automated analysis method to reduce operator dependence, correct for B_1 inhomogeneity and vascular structures, and extend the analysis to display multiple intensity clusters.

Materials and Methods—Two segmentation methods were compared—seeded region-growing, previously validated by expert reader scoring, and a new linear-binning method that corrects for the effects of bias field and vascular structures. The new method removes nearly all operator intervention by re-scaling the ^{129}Xe MRI to the 99th percentile of the cumulative distribution and applying fixed thresholds to classify ^{129}Xe voxels into 4 clusters: defect, low, medium, and high intensity. The methods were applied to 24 subjects including chronic obstructive pulmonary disease (COPD) subjects (n = 8), age-matched controls (n = 8), and healthy normal subjects (n = 8).

Results—Linear-binning enabled a faster and more reproducible workflow, and enabled analysis of an additional 0.25 ± 0.18 L of lung volume by accounting for vasculature. Like region-growing, linear-binning VDP correlated strongly with reader scoring (R^2 = 0.93, $p<0.0001$), but with less systematic bias. Moreover, linear-binning maps clearly depict regions of low and high intensity that may prove useful for phenotyping subjects with COPD.

Corresponding Author: Bastiaan Driehuys, Center for In Vivo Microscopy, Box 3302, Duke University Medical Center, Durham, NC 27710, Phone (919) 684-7786, Fax (919) 684-7158, bastiaan.driehuys@duke.edu.

Conclusions—Corrected linear-binning provides a robust means to quantify ^{129}Xe ventilation images, yielding VDP values that are indistinguishable from expert reader scores, while exploiting the entire dynamic range to depict multiple image clusters.

Keywords

Hyperpolarized ^{129}Xe MRI; ventilation defect analysis; lung segmentation; image registration

INTRODUCTION

The introduction of hyperpolarized (HP) ^{129}Xe MRI into clinical research has accelerated in recent years, with demonstrations of high-resolution imaging of pulmonary ventilation (1), alveolar microstructure (2), and gas exchange (3, 4). Like ^{3}He MRI, which emerged over the past decade, but suffers from supply limitations (5), ^{129}Xe MRI offers a means to enable radiation-free longitudinal imaging of pulmonary function. While recent interest in ^{129}Xe has been centered on exploiting its intriguing properties such as solubility and chemical-shift to probe diffusive gasexchange (2), important questions still remain to be addressed about how best to acquire and analyze even basic ^{129}Xe spin-density (ventilation) images. HP ^{129}Xe spin-density imaging, like ^{3}He MRI, readily depicts ventilation impairment in a variety of pulmonary disorders such as asthma, chronic obstructive pulmonary disease (COPD), cystic fibrosis (CF), and radiation-induced lung injury (RILI) (6–9).

While much progress has been made using reader-based scoring of HP gas MRI, it is also important to quantify these images in a way that is robust and automated. Such quantification not only enables high-throughput imaging, but also provides a means to extract the full richness of functional lung imaging. The simplest, and longest-standing quantification metric is the ventilation defect percentage (VDP), introduced by Woodhouse et al (10). Methods to calculate VDP were subsequently extended by Kirby et al. (9) who exploited a breath-hold ^{1}H anatomical image to confine the analysis to the thoracic cavity. We have previously employed this basic framework in conjunction with supervised region-growing methods to quantify VDP of ^{129}Xe MRI (6). This approach showed good inter-operator agreement, and correlated well with expert reader scores. However, the study also highlighted the need to further remove operator dependence, and to extend the range of analysis beyond simply calculating defect percentages. Moreover, the automated analysis can also be confounded by two additional factors. Firstly, the flexible vest coil may produce an inhomogeneous B_1 field, which causes a bias field variation across the lung (11). Secondly, the presence of pulmonary vasculature within the thoracic cavity can cause ^{129}Xe signal intensity voids (12) that can be incorrectly classified as ventilation defects.

Thus, in this work we introduce several extensions of ^{129}Xe MR ventilation analysis and incorporate several emerging ideas from the ^{3}He literature. First, we correct for the non-uniform B_1 field using a bias-field correction (11). Second, we remove the influence of the pulmonary vasculature in our thoracic cavity mask using a ‘vesselness’ filter (12). Lastly, we introduce a robust means of re-scaling the image histogram to enable using the entire dynamic range of ventilation information, and cast this information into four different clusters. These extensions decrease processing time, reduce operator dependence, and

improve the robustness and reproducibility of the method, while improving ease of visualization, and providing a richer set of quantitative metrics for analysis.

MATERIALS AND METHODS

All studies were approved by Institutional Review Board and prior to enrollment, written informed consent was obtained from all subjects. Image analysis was conducted using data acquired during a previously reported clinical trial of hyperpolarized ^{129}Xe (6), and hence the acquisition methods are only briefly summarized here. To test the new analysis methods, a total of 24 ^{129}Xe ventilation images were analyzed. This included 8 younger healthy volunteers (HV; 6 female, 2 male; mean age 39.1 ± 13.6 years); 8 subjects with COPD (2 female, 6 male; mean age 66.9 ± 4.9 years; 7 with GOLD (13) stage II and 1 with GOLD stage III); and 8 age-matched controls (AMC; 5 female, 3 male; mean age 62.0 ± 8.3 years). All subjects with COPD had a smoking history of at least 20 pack-years, while AMC subjects had less than a 5 pack-year history and had not smoked in the 5 years prior to MRI.

^{129}Xe Polarization and Delivery

Isotopically enriched ^{129}Xe gas (85% ^{129}Xe, Spectra Gases Inc., Alpha, NJ) was polarized to 6–10% by rubidium vapor spin-exchange optical pumping and cryogenically accumulated (14) using a commercial polarizer (Polarean Inc., Model 9800; Durham, NC). HP ^{129}Xe was then thawed in 1-L volumes into a Tedlar bag (Jensen Inert Products, Coral Springs FL), and the polarization was determined using a polarization measurement station (Polarean Inc., Model 2881).

Image Acquisition

All the MR scans were performed on a 1.5 T GE Healthcare EXCITE 15M4 MR system. Subjects were fitted with a flexible chest coil (Clinical MR solutions, Brookfield, WI) that was tuned to the ^{129}Xe Larmor frequency 17.66 MHz and proton-blocked to permit anatomical scans to be acquired using the ^{1}H body coil. After the initial localizers and proton thoracic cavity scans (described below), subjects underwent ^{129}Xe ventilation MRI. All scans were performed in the supine position. Prior to imaging, all subjects practiced achieving reproducible breath-holds by inhaling 1-liter of air from a bag starting from functional residual capacity (FRC). Prior to ^{129}Xe MRI, anatomic ^{1}H images were acquired during a 16-second breath-hold. These images were acquired using a steady-state free precession (SSFP) sequence (FOV = 40×40 cm^2, 128×128 matrix, 15-mm slices, TR/TE = 2.8/1.2 ms, α = 45°, BW = 125 kHz). This sequence highlights the pulmonary vasculature, which is useful for generating accurate thoracic cavity masks for quantifying ventilation. HP ^{129}Xe ventilation images were acquired during a separate 8.5–12 second breath-hold, with 12 slices acquired in the anterior-to-posterior direction using a spoiled gradient recalled echo (SPGRE) sequence (FOV = 40×28–40 cm^2, matrix = 128×90–128, 15-mm slices, TR/TE = 7.9/1.9 ms, α = 5–7°, BW = 8.0 kHz).

Semi-automated Image Analysis Methods

Creating the Registered Thoracic Cavity Mask—All image processing used a registered ^{1}H image of the thoracic cavity to develop a mask within which the ^{129}Xe

ventilation image analysis was constrained. ^{1}H thoracic cavity images were first registered to the ^{129}Xe images using a multi-resolution affine transform (15). All registration steps were completed using the Image Registration Toolkit (Visual Information Processing Group, London, UK) (16). Registration was followed by semi-automated segmentation of the ^{1}H images to generate a thoracic cavity mask. Thoracic cavity segmentation was performed using Avizo (Visualization Sciences Group, Burlington, MA), which employed the seeded region-growing algorithm previously described (6). Briefly, this involved the user placing a seed voxel in the lowest intensity of the thoracic cavity's center slice. The seeded region-growing algorithm then constructed a volumetric segmentation of the thoracic cavity by labeling all voxels connected to the seed voxel that satisfied a user-adjusted intensity threshold. The resulting mask was then morphologically filled in MATLAB (MathWorks Inc., Natick, MA) to remove holes from the thoracic cavity mask (6).

Method 1: Seeded Region-growing Segmentation to Determine VDP—The first method employed to calculate VDP was seeded region-growing, which was described previously (6) and is depicted schematically in Figure 1. Just as for ^{1}H segmentation, region-growing was also used to segment the ^{129}Xe ventilation image and generate a ventilation mask. This provided a binary ventilation mask from which the major airways were removed via multiplication by the thoracic cavity mask. Ventilation defects were then defined as those regions within the thoracic cavity mask that were not found in the ventilation mask. The volume of these defects was then divided by the thoracic cavity volume to calculate the ventilation defect percentage (VDP).

Method 2: Scaled Linear-Binning Maps with Bias Field and Vesselness Corrections represent—As suggested by Virgincar et al. (6), once the thoracic cavity mask has been co-registered to the ^{129}Xe ventilation image, a far richer, and potentially more robust analysis of the entire ventilation distribution becomes possible. Thus, rather than segmenting the image into two clusters (defect and ventilated), the ^{129}Xe distribution within the thoracic cavity can be quantified automatically into four bins: signal void (1st cluster: defect), low (2nd cluster), medium (3rd cluster) and high intensity (4th cluster). However, prior to this classification, the images were subjected to two additional corrections: application of a 'vesselness' filter to better account for the exclusion of ^{129}Xe signal by the pulmonary vasculature, and a bias-field correction to account for coil B_1 inhomogeneity. The associated workflow to generate such a corrected linear-binning map is depicted in Figure 2.

Linear-binning Histogram Scaling: Like the seeded region-growing method, the linear-binning method also relies on classification of the intensity histogram. The native histogram typically contains a small number of pixels that represent a high-intensity 'tail' that can confound the classification. However, these high-intensity voxels, which likely represent signal from airways that will contain 100% ^{129}Xe (17) and can therefore serve as a useful reference standard against which to scale the histogram. To consolidate this tail and establish a firm upper intensity cutoff, a threshold was set at the 99th percentile of the cumulative distribution. This was then used to equalize the histogram, and rescale it to range from 0 to 1. After rescaling, pre-determined thresholds were used to assign each ^{129}Xe voxel

to a particular intensity cluster. The thresholds for the 4 intensity clusters are defined as follows: defect (0–0.2), low (0.2–0.4), medium (0.4–0.8) and high intensity (0.8–1).

‘Vesselness’ Correction of the Thoracic Cavity Mask: The larger pulmonary vasculature creates regions that naturally exclude ^{129}Xe signal. Unless this vasculature is accounted for when defining the thoracic cavity mask, it will contribute to a false increase in the VDP. Previously, this problem was addressed by Tustison et al. (18), who segmented the vasculature directly from HP ^{3}He ventilation images. However, given that HP ^{129}Xe MRI exhibits somewhat lower SNR and spatial resolution than ^{3}He, we elected to extract the vasculature from the proton image (Figure 3a). After segmenting the thoracic cavity mask from the proton image, the holes inside the mask were morphologically closed (Figure 3b). Next, the vasculature was segmented from the thoracic cavity image by applying the vessel-enhancing algorithm first introduced by Frangi et al. (12), using a routine written in MATLAB. This approach exploits three properties of the vessel-like structure: intensity difference, tubular shape, and tree-like structure to indicate bifurcation and connectedness. This ‘vesselness’ filtering algorithm begins by calculating the Hessian matrix of a given image convoluted with a Gaussian kernel applied at a range of manually adjusted spatial scales. At each spatial scale σ, tubular structures are detected within the 3D volume by calculating three eigenvalues from the Hessian matrix. In a vessel, one of these values will be small (λ_1), while the other two (λ_2, λ_3) are large and of equal sign. These conditions indicate that a particular pixel belongs to the vessel tree under consideration. Meanwhile, the eigenvector u_1 defines the direction along the vessel while the other two eigenvectors (u_2, u_3) form the orthogonal plane of the vessel. In this way, different vessel sizes will be detected at their corresponding scales. Therefore, the vasculature was detected and removed to establish a corrected mask (Figure 3d) that preserves the maximum possible ^{129}Xe image volume for quantitative analysis.

Retrospective Bias-Field Correction: Because our images were acquired with a flexible vest coil, the resulting RF field across the chest cavity is not completely homogeneous. This results in application of slightly varying flip angles and coil sensitivities across the field of view, and manifests itself as a slowly varying inhomogeneity or bias-field (19, 20). Here, we incorporate the N4ITK retrospective bias-field correction in ANTS (Advanced Normalization Tools; http://picsl.upenn.edu/ANTS/), which Tustison et al. first introduced and applied to HP ^{3}He MRI (11, 20). This method assumes a spatially smooth and slowly varying intensity bias across the image, and estimates it using a B-spline approximation, which is iteratively adapted to correct the resulting images.

Statistical Methods

All statistical analysis was performed using JMP 11 (SAS Institute Inc., Cary, NC). The relationship between VDP calculated by seeded region-growing, linear-binning, and expert reader scoring (6) was determined using linear regression analysis and Pearson correlation coefficients (r). A Bland-Altman analysis was employed to test for systematic differences between the segmentation and quantification methods. To evaluate the ability of the various segmentation methods to reveal significant differences between groups, the one-way

analysis of variance (ANOVA) test was employed. For all correlations and comparisons, the level of significance was 5% ($p<0.05$).

RESULTS

Linear-binning Histogram Scaling

Figure 4 illustrates the importance of correctly handling the high-intensity "tail" (Figure 4d) in the ^{129}Xe intensity histogram prior to rescaling to the range of 0–1. A previous approach we had suggested (6) was to divide all intensities by the average of the top 5% of values. As illustrated in Figure 4e, this scaling does reduce the tail, but does not completely remove it. The rescaled histogram is still significantly weighted towards low-intensity values and the resulting maps significantly overestimate the number of pixels belonging to the defect and low-intensity clusters, while underestimating medium-intensity clusters and exhibiting no high-intensity values (Figure 4b). By instead rescaling by the 99th percentile of the cumulative distribution, the high-intensity tail is effectively removed and the scaled histogram now encompasses the full range from 0–1, with the bulk of the distribution falling near the middle of the range (Figure 4f). The resulting binning maps now better reflect reader perception and highlight not just ventilation defects, but also yield a majority of medium intensity pixels, while appropriately visualizing the low and high intensity clusters.

'Vesselness' Filter

Figure 5 shows the different approaches for incorporating the effects of pulmonary vasculature to create the thoracic cavity mask. As shown in Figure 5a–b, this healthy volunteer exhibits homogeneous ventilation throughout the lung, except for those regions containing pulmonary vasculature. As shown in Figure 5c, the initial thoracic cavity mask created by seeded region-growing does exclude this vasculature from the mask but when this mask is applied to the ^{129}Xe ventilation image in Figure 5d, it clearly removes too much of the lung from the analysis. This problem can be somewhat diminished by using a morphological closing operation (6) as in Figure 5e, but not only does this still remove too much of the mask near the large vasculature, it also completely eliminates the smaller vessels from the mask. In doing so, the regions where ^{129}Xe signal is excluded by the smaller vessels, will be erroneously counted as ventilation defects. We estimate that the pulmonary vasculature contributes an additional volume of about 0.25 ± 0.18 L to the thoracic cavity, and if not accounted for, will cause VDP to be overestimated by 2–6%. As shown in Figure 5g–h, when the 'vesselness' filter is applied, the vasculature is well represented in the refined mask. This serves to maximize the volume of lung that can be analyzed, while excluding regions of the lung mask where vasculature truly excludes ^{129}Xe signal.

Bias-field Correction

Figure 6 shows the effects of correcting for the bias field imposed by the high-sensitivity, but flexible vest coil. Note in subject 1, the high signal intensity appearing in the lower left lung. This high-intensity region appears to be caused by a small patch of very high coil sensitivity that was 2.5× higher than the mean. Upon removing this bias field, this region now matches that of its surroundings. Overall, this subject showed a range of bias fields that

contributed to signal intensities that were between −20% and +150% above the mean. However, the bias-field ranged from a more modest −20% to +20% over most of the image. As illustrated by subject 2, the strongest bias-fields were observed near the anterior and posterior chest walls, close to the 4 coil elements in this flexible coil.

Seeded Region-growing and Linear-binning Clustering Measurements

All images (n = 24) that had previously undergone seeded region-growing segmentation (6) were also successfully segmented by the corrected linear-binning clustering methods described here. Figure 7 shows three sample datasets from a 35-yr-old healthy volunteer, a 65-yr-old age-matched control, and a 73-yr-old COPD subject. The healthy volunteer exhibits homogeneous ventilation in most lung regions, with the exception of lower signal intensities seen in the apex and peripheral regions of the lung. This subject shows very few low-intensity or high-intensity clusters, and thus seeded region-growing and linear-binning agree well. The analysis becomes more nuanced with the age-matched control. Again, both methods appear to accurately depict the clearly visible small ventilation defects, primarily located in the periphery of the lung. However, in this case, linear-binning now also depicts significant areas of low-intensity (19.9% by volume in this subject), as well as some high-intensity (6.9% by volume), primarily in periphery and apex of the lungs. Finally, in the COPD subject, linear-binning identifies a significantly larger VDP than region-growing. The higher VDP is actually in closer accordance with expert scoring. As with the age-matched control subject, linear-binning also identifies significant regions of both low and high intensity.

As shown in Figure 8, all three methods of image quantification (VDS%, region-growing VDP, and linear-binning VDP), showed significant differences in ventilation defect levels between the three groups ($p<0.0001$, one-way Anova; see Table 1). All three methods separated AMC from COPD ($p<0.0001$), as well as HV from COPD ($p<0.0001$), but none differentiated AMC from HV in this small sample. In HVs, reader scoring showed the greatest dispersion across subjects, while region-growing provided tightly clustered, low VDP values. Compared to region-growing, linear-binning generated a slightly higher mean VDP in the HVs, where subtle defects predominate. In AMCs, all three methods reported a generally higher VDP, and the differences between methods were similar to what was seen in HVs. Region-growing reported lower VDP and less variance than expert reader scoring, while linear-binning appeared to split the difference. The differences between methods were most apparent in the COPD group, where region-growing generated a lower VDP than the other methods, while linear-binning results were more closely aligned with expert reader scoring.

Unlike reader scoring and region-growing, which provide only defect scores and percentages, linear-binning also reveals and reports on the distribution of low, medium and high intensity regions. The high-intensity cluster was significantly different between HVs and COPD subjects, but did not separate HVs from AMCs, nor AMCs from COPD subjects. The medium cluster was significantly different between all groups ($p<0.0001$, one-way Anova), with COPD subjects having a medium-intensity cluster fraction of only 32.7 ± 14.2% (range: 14.0–48.6%), compared to 65.9 ± 11.8% for AMCs ($p<0.0001$) and 70.9 ±

4.7% for HVs ($p<0.0001$). Interestingly, the percentage of high-intensity voxels was greatest for the HVs for whom it comprised 8.73 ± 1.28% of their volume, compared to 7.90 ± 3.36% for AMCs ($p=0.0017$) and only 3.92 ± 1.37% for COPD ($p=0.0003$).

Correlations and Bland-Altman Analysis

Figure 9 shows correlation and Bland-Altman plots comparing seeded region-growing and linear-binning-derived ventilation defect percentage (VDP) to expert reader ventilation defect scores (VDS%). Both region-growing and binning correlated significantly with reader scores. For region-growing the correlation to expert readers was $R^2 = 0.928$ ($p<0.0001$), while linear-binning provided a similar correlation ($R^2 = 0.934$). However, Bland-Altman analysis shows that region-growing provided a consistently lower defect percentage than readers, with biases ranging from −5%–30%, and trending towards greater differences at the highest defect percentages. Meanwhile, a similar Bland-Altman analysis comparing linear-binning-derived VDP to expert reader VDS%, still shows some significant biases between the methods, but there is no longer a clear trend. Moreover, biases between expert reader VDS% and linear-binning VDP were considerably smaller and ranged from only −10%–15%.

DISCUSSION

Advantages of the Corrected Linear-binning Method

The corrected linear-binning method provides an elegant way to quantify and visualize ventilation defects and has the capability to highlight subtle features of ventilation distribution that could be missed by the simple binary classification. Moreover, linear-binning mitigates the subjective elements of the analysis caused by intra- and inter-observer bias while providing a substantial throughput improvement. The only user intervention in the process is initialization and supervision of the ^{1}H segmentation. Moreover, this workflow is naturally amenable to applying corrections for B_1 inhomogeneity associated with flexible coils. Although the importance of correcting for such non-physiological sources of hyper- or hypointensity was noted by Tustison et al (18), it is not yet widely practiced in the HP gas MRI analysis. Moreover, we have further refined the analysis by more carefully accounting for the role of vasculature in the thoracic cavity, thereby reducing false VDP by an estimated 2–6%. In aggregate, this analysis method should significantly expand our ability to mine information from ^{129}Xe ventilation images.

Comparison of Reader-based Scoring, Seeded Region-growing and Linear-binning Methods

As was previously noted for the case of seeded region-growing, VDP calculated by corrected linear-binning also correlated exceptionally well with expert reader VDS% ($R^2 = 0.934$, $p<0.0001$). However, a key difference between the two segmentation methods is revealed by the Bland-Altman analysis showing that region-growing VDP exhibits a systematic bias towards lower defect percentages relative to reader VDS%. This is likely because region-growing segments the ventilation image without including thoracic cavity mask information. In this case some background noise voxels in the ventilation image may share the same intensity as some less-ventilated regions within the lung. This requires the

user to select a somewhat lower threshold in order to segment only those voxels perceived to lie within the lung, without inadvertently 'growing' to the background. This has the effect of excluding some poorly ventilated voxels in the lung periphery. Of course, linear-binning does not suffer the same problem, because the analysis is already confined to the thoracic cavity mask and does not require region-growing. This may explain why linear-binning and associated re-scaling of the histogram generate VDP values that more closely resemble the interpretation of expert readers who make this calculation intuitively. Thus, linear-binning not only correlated well with expert reader VDS%, but also provided defect percentages of similar magnitude.

It is interesting to single out and more carefully consider some cases in which reader VDS% and linear-binning VDP did not agree perfectly. As seen in Figure 8, there were several subjects who received a reader score with VDS% = 0%, while linear-binning identified a non-zero VDP. It is certainly possible that the binning method could over-estimate ventilation defects when the thoracic cavity and ^{129}Xe images are misregistered. In a small fraction of cases we noted a significantly different diaphragm position during the two breath-holds, and these were more challenging to register. However, it is also conceivable, as suggested by Virgincar et al., (6), that readers tend to ignore ventilation defects if they comprise less than 2% of the total volume, whereas the linear-binning method and seeded region-growing method consider every voxel within the thoracic cavity independently, regardless of the context. But perhaps a more interesting interpretation is provided in Figure 10, where the top row shows a healthy volunteer who received a score of VDS%=0. Indeed, based on viewing purely the grey-scale information, this image appears fully homogeneous. However, when the registered thoracic cavity is also included, the linear-binning technique makes apparent a thin rind of ventilation defect (VDP=8.44%) and low intensity regions (13.50%) primarily in the apex of the right lung. Hence, in this situation, the additional context provided by the thoracic cavity reveals ventilation defects and low-intensities that are not readily apparent when viewing the gray-scale image in isolation.

A second example of discordance between expert reader and linear-binning is shown in the bottom of Figure 10. In this case, the reader VDS% is 16.67%, whereas binning assigned a lower VDP of 6.65%. In fact, this 74-yr-old AMC had a somewhat tortuous thoracic cavity that resulted in a jagged looking ventilation image. However, when this thoracic cavity is incorporated into the VDP calculated by binning, the defect percentage is quite modest, and significantly lower than the reader score. These two examples, thus illustrate that accurate incorporation of the thoracic cavity mask is critical for both avoiding under- and overestimation of ventilation defects.

Having established that the linear-binning VDP correlates well with expert reader scoring, and that minor differences between the two, may be more attributable to slight inaccuracies in the way readers estimate the thoracic cavity, it is useful to consider the value of clusters beyond the ventilation defect percentage. For example, in HVs and AMCs, there is a clear tendency towards detecting low-intensity signal in the peripheral lung. Such regions may be associated with aging-related airway narrowing or loss of compliance. Similarly, the corrected binning reveals significant regions of high-intensity signal in the HV and AMC. The physiological interpretation of these clusters is not yet apparent, but clearly provides

another possible metric for study. Note that we have elected to consider only 4 levels of intensity in this analysis; the method can be easily extended to any number. The utility of this extension will have to be determined by scanning greater numbers of subjects with wider ranges of disease. However, it is conceivable that careful analysis of these additional clusters would provide insights into different COPD phenotypes or be predictive of exacerbations.

Study Limitations

While the results of this analysis method are encouraging, the study did have some limitations. First, since the thoracic cavity and ventilation images were obtained during different breath-holds, body movement, and lung volume differences could create difficulties in registering the two images. Thus, there may be benefit in simultaneously acquiring the ^{1}H thoracic cavity image and the HP gas ventilation image within a single breath-hold, as described by Wild et al. (21). Such an approach should largely obviate the need for registration, or at least greatly reduce the degree of image warping required. Prior to implementing such measures, we have recently found that diaphragm position is significantly more reproducible between scans if subjects inhale air from a 1-liter bag during the ^{1}H thoracic cavity image rather than inhaling freely.

Further, it should be noted that the binning thresholds of 0.2, 0.4, and 0.8 used to divide the scaled ventilation intensity histograms into clusters were based on qualitative visual assessment. It is conceivable that further refinement of these thresholds could hold further benefit. Furthermore, it should be noted that our cluster segmentation was based entirely on intensities and did not exploit spatial contextual information such as k-means cluster analysis (9). Future work may incorporate such information as done, for example, in the Atropos software package introduced by Tustison et al. (18). Finally, the bias-field correction algorithm used to correct B_1 inhomogeneity is derived primarily from the brain MRI literature and the method was modified by Tustison et al. (22) to enable correction of ^{3}He ventilation images. However, further work is needed to ensure that such corrections do not inadvertently remove physiologically important contributions to intensity non-uniformity, such as posture-related ventilation gradients (2). We intend to pursue phantom studies to further investigate and document the validity of such bias field corrections in a simple system with no known physiological gradients.

CONCLUSIONS

This study demonstrated the feasibility of using corrected linear-binning analysis of ^{129}Xe MRI to visualize and quantify regional ventilation in healthy volunteers, age-matched controls and COPD patients. The resulting defect percentages correlated exceptionally well and were of similar magnitude to expert reader scores, suggesting that linear-binning closely follows expert reader observations. Moreover, this linear-binning method provides a promising means to not only accelerate objective image analysis, but by straightforward extension can provide other measures of regional function such as regional distribution and heterogeneity. While demonstrated here for patients with COPD, this quantitative technique

is likely to be equally applicable to visualization and quantification of ^{129}Xe MRI in a broad variety of pulmonary diseases.

Acknowledgments

This study was funded NIH/NHLBI R01HL105643 with additional support from GE Healthcare. Analysis was conducted with additional support from the Duke Center for In Vivo Microscopy, a NIH/NIBIB national Biomedical Technology Resource Center (P41 EB015897). The authors also want to thank Sally Zimney for careful proofreading of the manuscript.

Grant Support:

NIH/NHLBI R01 HL105643, NIBIB P41 EB015897

References

1. Svenningsen S, Kirby M, Starr D, et al. Hyperpolarized He and Xe MRI: Differences in asthma before bronchodilation. J Magn Reson Imaging. 2013; 38(6):1521–1530. [PubMed: 23589465]
2. Kaushik SS, Freeman MS, Cleveland ZI, et al. Probing the Regional Distribution of Pulmonary Gas-Exchange through Single-Breath, Gas- and Dissolved-Phase 129Xe MR Imaging. J Appl Physiol. 2013; 115(6):850–860. [PubMed: 23845983]
3. Qing K, Ruppert K, Jiang Y, et al. Regional mapping of gas uptake by blood and tissue in the human lung using hyperpolarized xenon-129 MRI. J Magn Reson Imaging. 2014; 39(2):346–359. [PubMed: 23681559]
4. Cleveland ZI, Cofer GP, Metz G, et al. Hyperpolarized Xe MR imaging of alveolar gas uptake in humans. PLoS One. 2010; 5(8):e12192. [PubMed: 20808950]
5. Dana AS, Morgan D. The Helium-3 Shortage: Supply, Demand, and Options for Congress. Congressional Research Service. 2010
6. Virgincar RS, Cleveland ZI, Sivaram Kaushik S, et al. Quantitative analysis of hyperpolarized (129) Xe ventilation imaging in healthy volunteers and subjects with chronic obstructive pulmonary disease. NMR Biomed. 2012; 26(4):424–435. [PubMed: 23065808]
7. Shukla Y, Wheatley A, Kirby M, et al. Hyperpolarized 129Xe magnetic resonance imaging: tolerability in healthy volunteers and subjects with pulmonary disease. Acad Radiol. 2012; 19(8): 941–951. [PubMed: 22591724]
8. Kirby M, Svenningsen S, Owrangi A, et al. Hyperpolarized 3He and 129Xe MR imaging in healthy volunteers and patients with chronic obstructive pulmonary disease. Radiology. 2012; 265(2):600–610. [PubMed: 22952383]
9. Kirby M, Heydarian M, Svenningsen S, et al. Hyperpolarized 3He magnetic resonance functional imaging semiautomated segmentation. Acad Radiol. 2012; 19(2):141–152. [PubMed: 22104288]
10. Woodhouse N, Wild JM, Paley MN, et al. Combined helium-3/proton magnetic resonance imaging measurement of ventilated lung volumes in smokers compared to never-smokers. J Magn Reson Imaging. 2005; 21(4):365–369. [PubMed: 15779032]
11. Tustison NJ, Avants BB, Cook PA, et al. N4ITK: improved N3 bias correction. IEEE Trans Med Imaging. 2010; 29(6):1310–1320. [PubMed: 20378467]
12. Frangi AF, Niessen WJ, Hoogeveen RM, van Walsum T, Viergever MA. Model-based quantitation of 3-D magnetic resonance angiographic images. IEEE Trans Med Imaging. 1999; 18(10):946–956. [PubMed: 10628954]
13. Rabe KF, Hurd S, Anzueto A, et al. Global strategy for the diagnosis, management, and prevention of chronic obstructive pulmonary disease – GOLD executive summary. American Journal of Respiratory and Critical Care Medicine. 2007; 176(6):532–555. [PubMed: 17507545]
14. Driehuys B, Cates GD, Miron E, Sauer K, Walter DK, Happer W. High-volume production of laser-polarized Xe-129. Applied Physics Letters. 1996; 69(12):1668–1670.
15. Studholme C, Hill DLG, Hawkes DJ. An overlap invariant entropy measure of 3D medical image alignment. Pattern Recogn. 1999; 32(1):71–86.

16. Rueckert D, Sonoda LI, Hayes C, Hill DLG, Leach MO, Hawkes DJ. Nonrigid registration using free-form deformations: Application to breast MR images. Ieee T Med Imaging. 1999; 18(8):712–721.

17. Mistry NN, Thomas A, Kaushik SS, Johnson GA, Driehuys B. Quantitative analysis of hyperpolarized 3He ventilation changes in mice challenged with methacholine. Magnetic resonance in medicine: official journal of the Society of Magnetic Resonance in Medicine/Society of Magnetic Resonance in Medicine. 2010; 63(3):658–666.

18. Tustison NJ, Avants BB, Flors L, et al. Ventilation-based segmentation of the lungs using hyperpolarized (3) He MRI. J Magn Reson Imaging. 2011; 34(4):831–841. [PubMed: 21837781]

19. Belaroussi B, Milles J, Carme S, Zhu YM, Benoit-Cattin H. Intensity non-uniformity correction in MRI: existing methods and their validation. Med Image Anal. 2006; 10(2):234–246. [PubMed: 16307900]

20. Arnold JB, Liow JS, Schaper KA, et al. Qualitative and quantitative evaluation of six algorithms for correcting intensity nonuniformity effects. Neuroimage. 2001; 13(5):931–943. [PubMed: 11304088]

21. Wild JM, Ajraoui S, Deppe MH, et al. Synchronous acquisition of hyperpolarised 3He and 1H MR images of the lungs – maximising mutual anatomical and functional information. NMR Biomed. 2011; 24(2):130–134. [PubMed: 20821726]

22. Tustison NJ, Altes TA, Song G, de Lange EE, Mugler JP 3rd, Gee JC. Feature analysis of hyperpolarized helium-3 pulmonary MRI: a study of asthmatics versus nonasthmatics. Magn Reson Med. 2010; 63(6):1448–1455. [PubMed: 20512846]

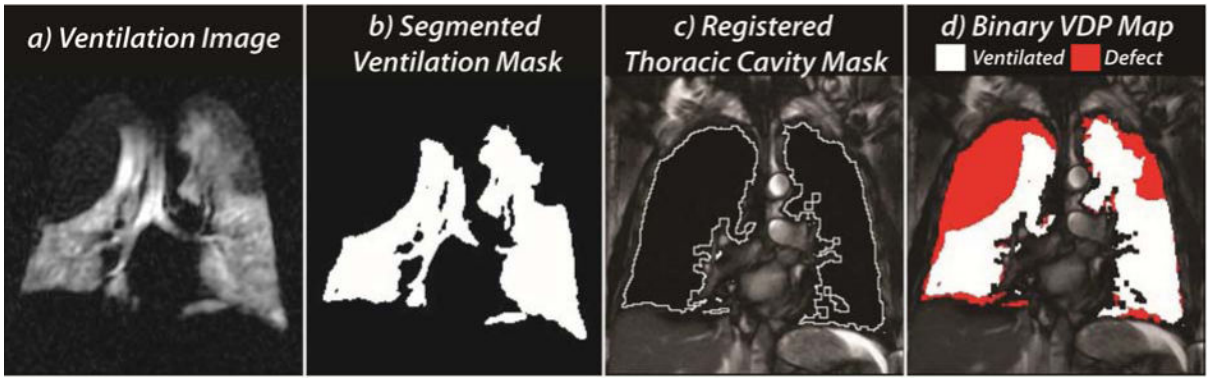

Figure 1.

Seeded region-growing segmentation workflow: the ^{129}Xe ventilation image (a) was segmented to obtain a ventilation mask (b). The ^{1}H image was registered to the ^{129}Xe image and segmented to obtain a thoracic cavity mask (c). These masks were multiplied to remove major airways from the ventilation mask, thereby obtaining the ventilated volume within the thoracic cavity, enabling calculation of the binary ventilation defect map (d) as well as ventilation defect percentage.

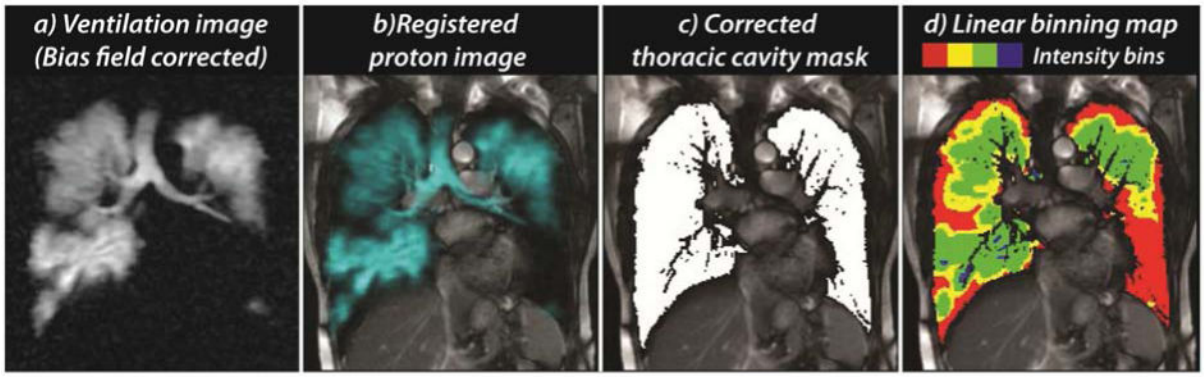

Figure 2.

Corrected linear-binning map workflow. The bias-field corrected ventilation images (a) were used to register the ^{1}H thoracic cavity images (b). These were then segmented and detected vasculature was removed to define the thoracic cavity mask that constrains the analysis (c). Pixels from the bias-field corrected ventilation image lying within the thoracic cavity volume were rescaled by their top percentile to range from 0-1 and classified into 4 intensity clusters to create the linear-binning map (d).

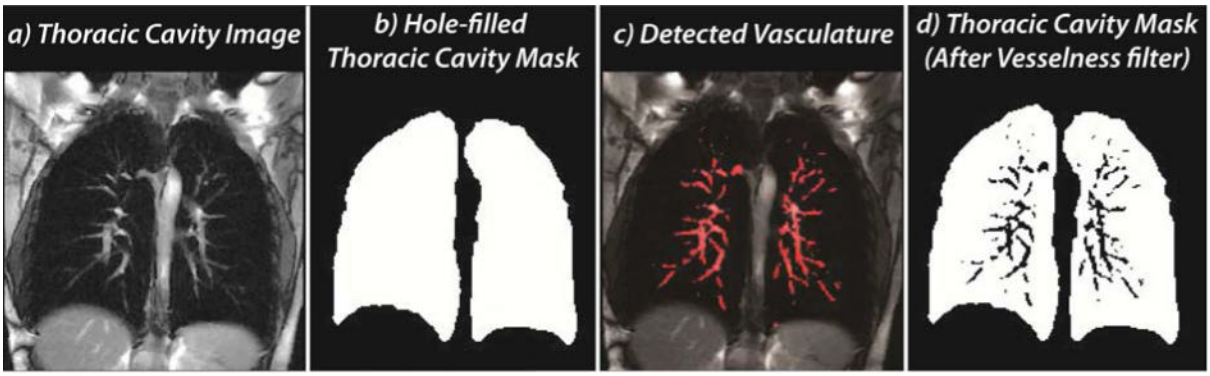

Figure 3.

Tuning of the thoracic cavity mask. The thoracic cavity image (a) was segmented and morphologically filled to generate the initial thoracic cavity mask (b). The thoracic cavity image then underwent vesselness filtering to detect the pulmonary vasculature (c). This was then removed from the original thoracic cavity image to generate the final thoracic cavity mask (d).

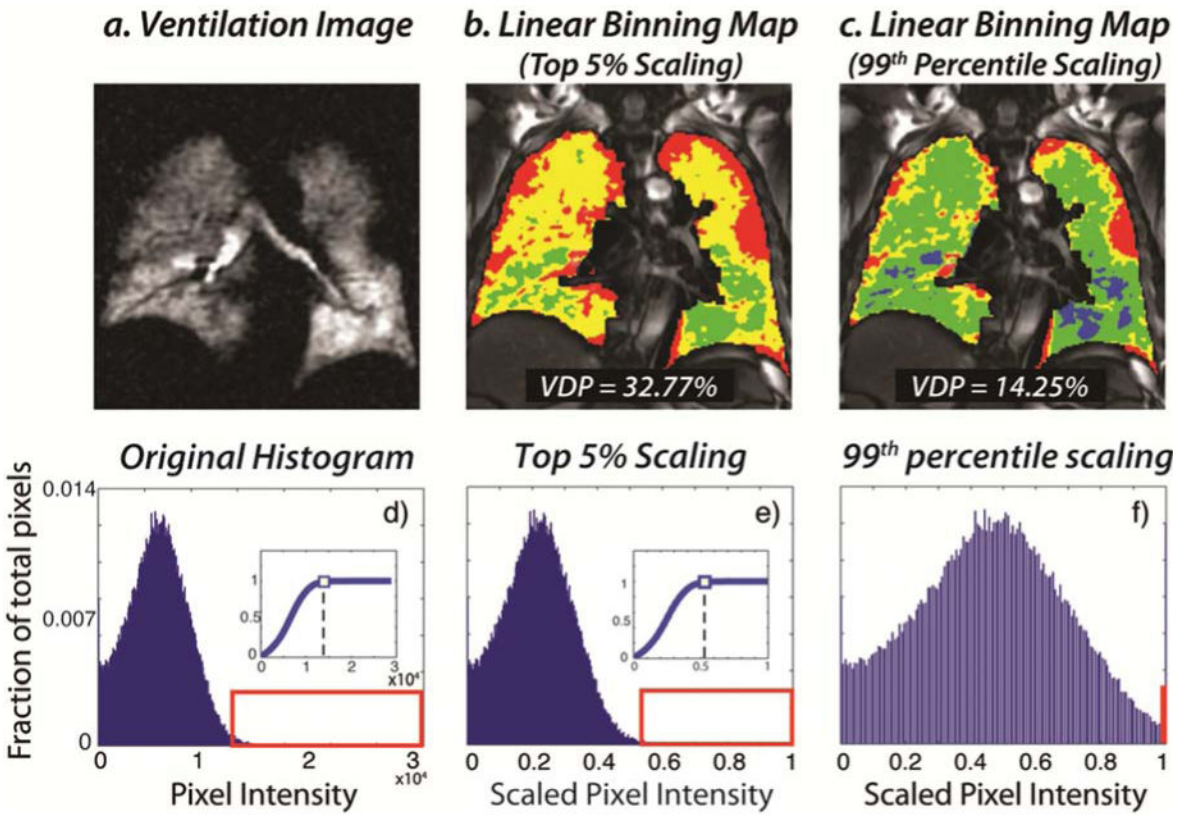

Figure 4.

^{129}Xe MRI histogram re-scaling: the native ^{129}Xe MR image (a) has a histogram with a high-intensity "tail" (d) that must be removed for effective rescaling. If this is done by simply dividing all intensities by their top 5%, the resulting binning maps (b) over-estimate the low-intensity bins because the tail has only been partially removed (e). By instead scaling all ^{129}Xe intensity pixels by the 99th percentile of the ^{129}Xe intensity cumulative distribution, the associated binning maps are more reproducible and more consistent with reader perception (c) because the histogram tail has been effectively removed (f).

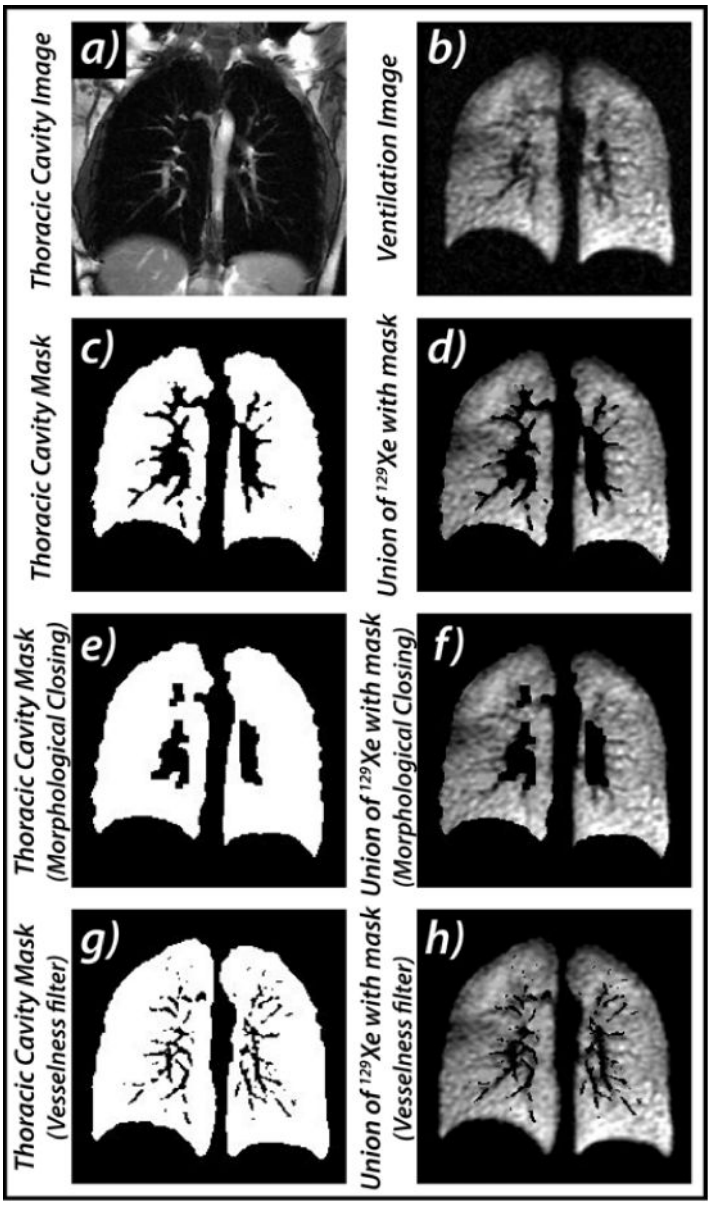

Figure 5.

Impact of the 'vesselness' filter. The thoracic cavity image (a) and ^{129}Xe ventilation image (b) in a healthy volunteer. The thoracic cavity mask after initial segmentation (c) overestimates the large vessels as shown in the union with ^{129}Xe MRI (d). The thoracic cavity mask after morphological closing (e) slightly reduces the over-estimation of the larger vessels, but completely eliminates the exclusion of smaller vasculature from the mask (f). The mask after application of 'vesselness' filter (g) now shows a thoracic cavity mask that best excludes the vasculature, while preserving a maximum ^{129}Xe image volume for quantitative analysis (h).

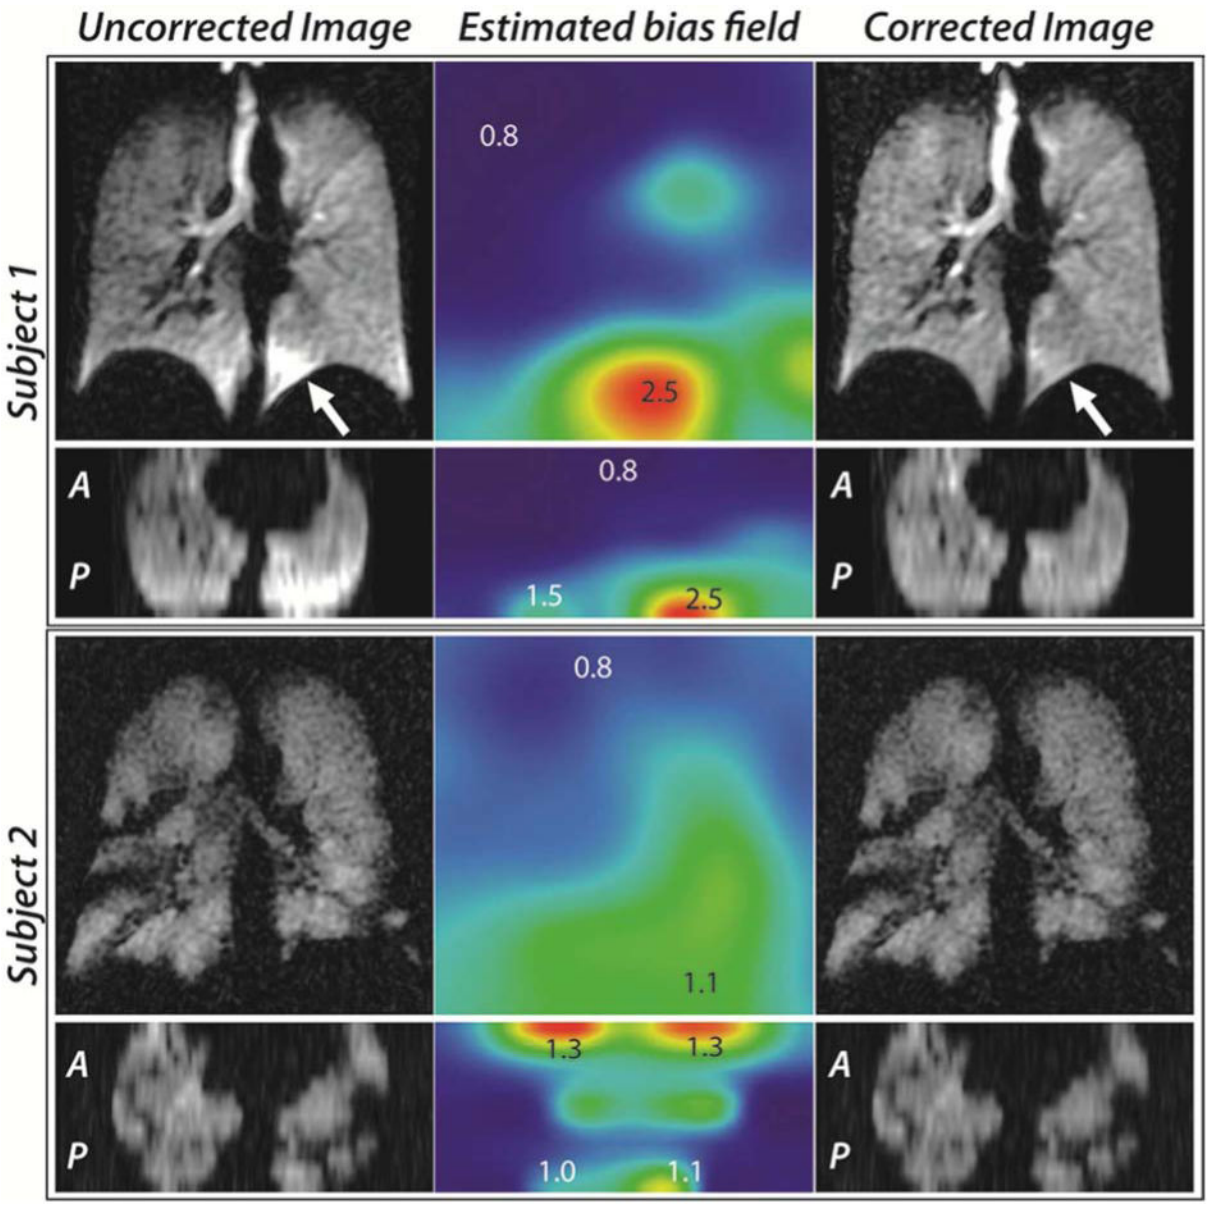

Figure 6.

Bias-field correction: Column 1 shows ventilation images from two subjects (coronal and axial views) prior to the application of the retrospective bias correction algorithm. Column 2 depicts the estimated bias-fields in both orientations (shown as a maximum intensity projection). Column 3 shows the ^{129}Xe images after bias-field correction. Bias fields appear most intense near the coil elements.

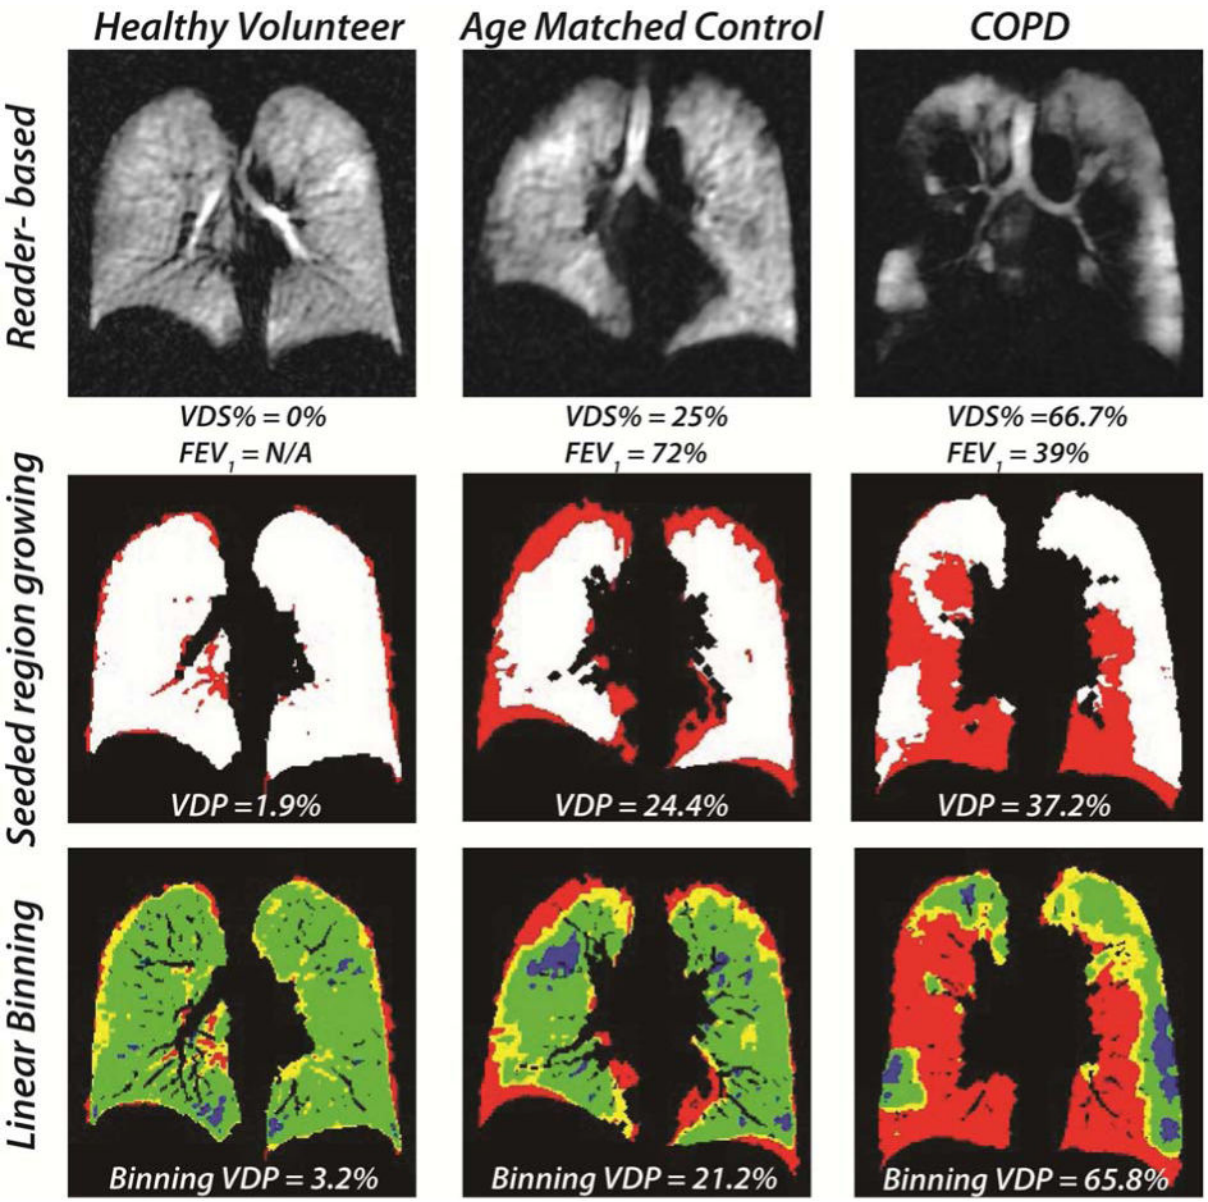

Figure 7.

Ventilation images (prior to bias-field correction) with associated reader-based scores, along with ventilation maps generated by the seeded region-growing and corrected linear-binning methods. Shown here are representative cases of a healthy volunteer, an age-matched control, and a COPD subject.

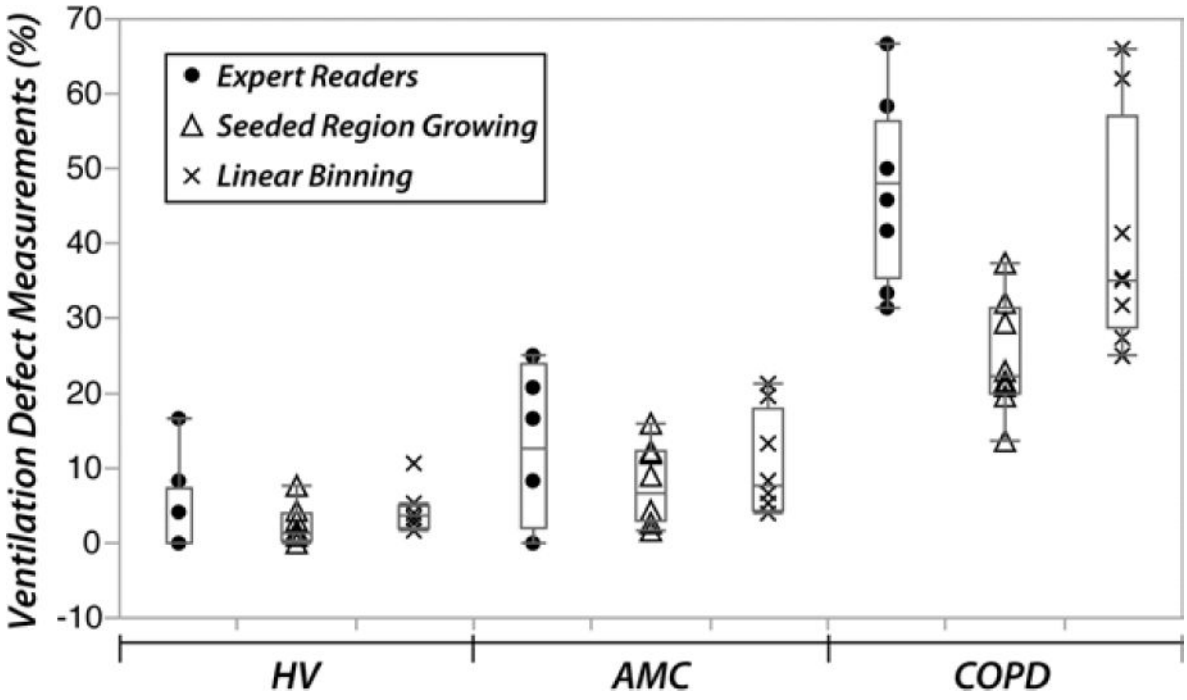

Figure 8.

Ventilation defect measurements for healthy volunteers, age-matched controls, and COPD subjects obtained by expert reader scoring, seeded region-growing, and linear-binning.

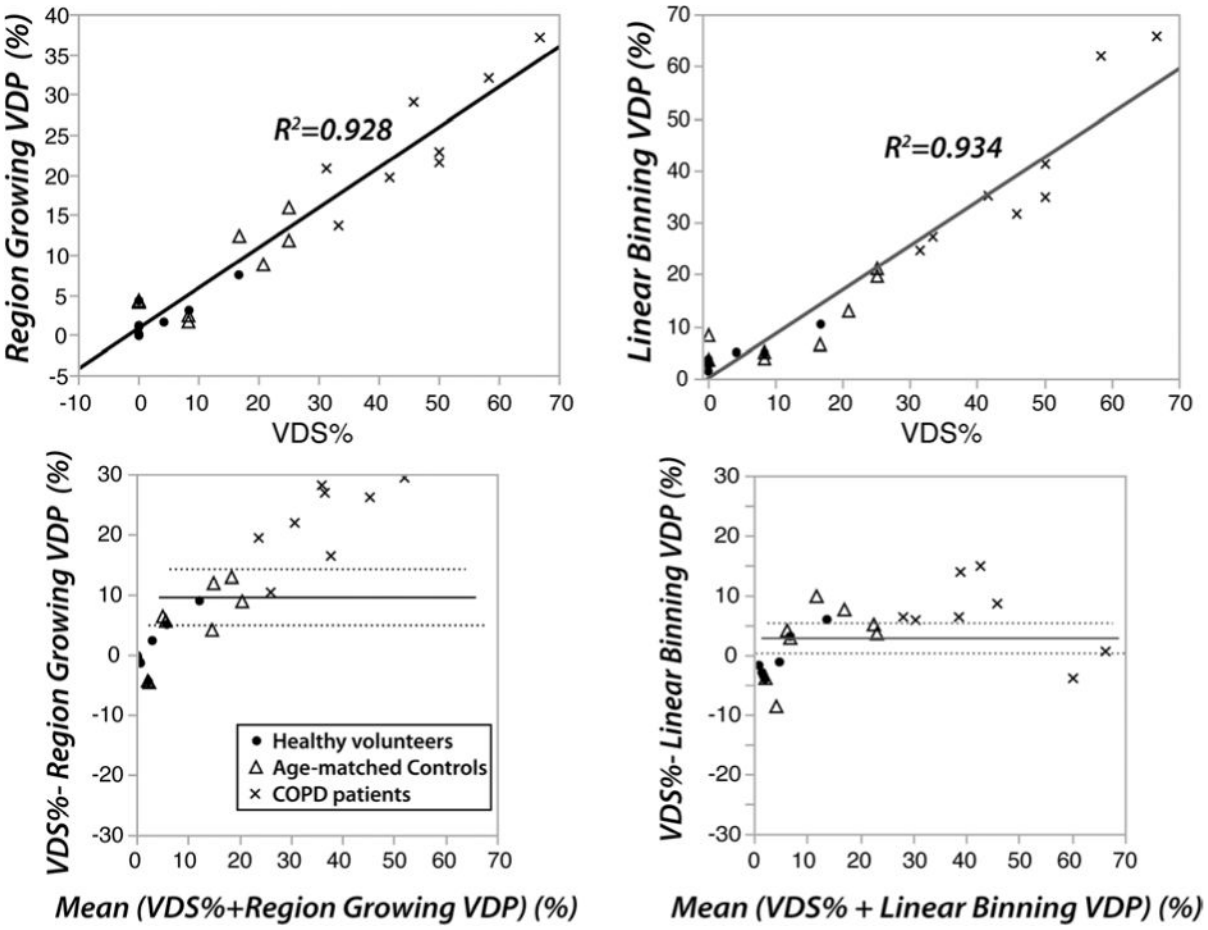

Figure 9.

Correlation of the expert reader-based VDS% with VDP calculated using the seeded region-growing and linear-binning methods. Note that although both methods correlate well with reader scores, Bland-Altman plots show a systematic bias in seeded region-growing that is significantly reduced for linear binning.

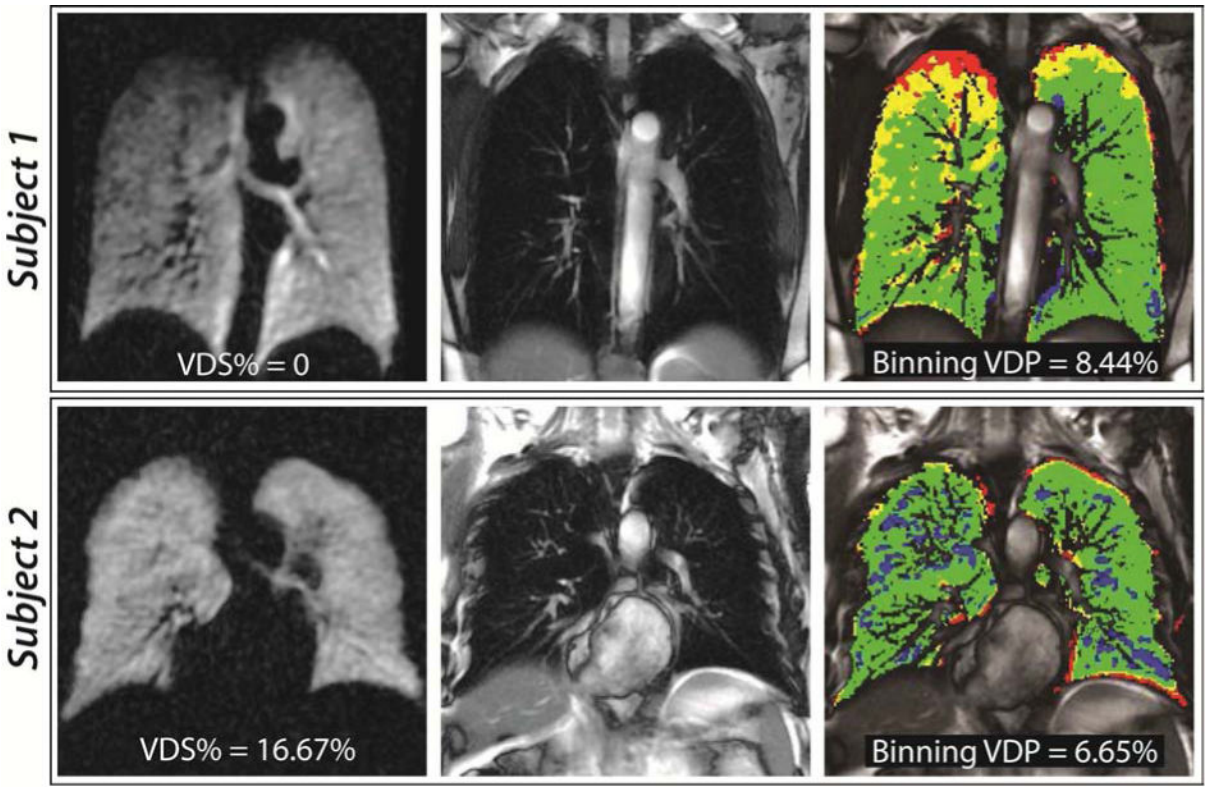

Figure 10.

Examples of slight discordance between reader scoring and binning. Subject 1 depicts a healthy volunteer appearing to have no clear defects, but when viewed in the context of the registered thoracic cavity, shows clear ventilation defects and low intensity primarily in the apex of the right lung. The second row shows images of an age-matched control with a somewhat tortuous thoracic cavity, which may have caused readers to assign a higher VDS % of 16.67%. However, binning analysis, which incorporates the tortuous thoracic cavity shows mostly medium intensity.

Table 1

Mean values (±standard deviation) of ventilation defect score percentage (VDS%) and ^{129}Xe cluster percentage measurements using region-growing and linear-binning

Parameter	HV (n = 8)	AMC (n = 8)	COPD (n = 8)	HV, AMC	AMC, COPD	HV, COPD
	Mean values ± standard deviation			P values		
VDS% (%)	3.65 ± 6.07	13.02 ± 10.31	47.14 ± 11.93	0.068	<.0001*	<.0001*
Seeded region growing						
VDP (%)	2.41 ± 2.55	7.77 ±5.27	24.68 ± 7.59	0.066	<.0001*	<.0001*
Linear binning method						
VDP (%)	4.24 ± 2.91	10.29 ± 6.96	40.45 ± 15.40	0.135	<.0001*	<.0001*
Hypo (%)	16.16 ± 3.76	15.94 ± 8.32	22.95 ± 7.07	0.947	0.055	0.048*
Normal (%)	70.87 ± 4.73	65.87 ± 11.76	32.69 ± 14.21	0.373	<.0001*	<.0001*
Hyper (%)	8.73 ± 1.28	7.90 ± 3.36	3.92 ± 1.37	0.463	0.0017*	0.0003*

HV: healthy volunteer; AMC: age-matched control; COPD: chronic obstructive pulmonary disease.

*Statistically significant