Chapter 4

Compensating PET: preliminary results

In this chapter the deformation vector fields output of the CT co-registration will be applied on PET images acquired in the same breathing phases. This should compensate for the blurring effect that the patient's respiratory motion has on the PET images. Three different tests will be performed: two of them based on a phantom thorax model, and another one based on the clinical PET/CT data provided by the Netherlands Cancer Institute.

A sample gated PET/CT data set from [28] was available, and makes it possible to evaluate the algorithm's performance with clinical data. The main disadvantage is that only one case was provided.

Another tool that has been used is the NCAT phantom [40], v1.13. This software phantom lets the user generate both CT and PET volumes from a simulated patient's thorax in different respiratory phases. It includes many configurable parameters, letting the user generate many different cases

4.1 Experiments with NCAT

4.1.1 Materials and methods

Materials

NCAT generates perfectly matched CT and PET volumes, as well as volumes containing only user-defined lesions (position and size; spherical shape is assumed). There are many different user-defined parameters, among them patient gender, diaphragm motion extension, body height or heart size. The images can hence be generated with a large number of degrees of freedom in the simulated anatomy. It also includes a lesion and a movement vector field generator, tools that will be very useful for evaluation purposes.

A virtual thorax was generated with NCAT's default parameters. The extent of the diaphragm and chest motions was however exaggerated (in a 15%, approximately) in order to make it easier to appreciate the respiratory motion's effect on the images. Ten different time frames equally distributed along the respiratory cycle were generated, as well as the corresponding movement vector fields from the first phase to the others.

16 volumes containing only a lesion with a diameter of 10 mm were also generated. 4 lesions were placed in each lung region: apical, central, base and peripheral

Methods

It would be initially feasible to directly compare the deformation fields estimated by the registration algorithms and the ones generated by the NCAT phantom tools. The closer these vectors are, the better. The problem with this method is that, as pointed out by the software author, the NCAT generated vector field has a constant value of zero in a large part of the volume, and is in general just an approximation of the "real" movements. Hence, this evaluation method will not be considered.

Another option to evaluate the registration algorithms with NCAT is to apply the estimated deformation fields (from the generated CT) on the volumes containing only the lesions. It is then possible to compare the lesion position in the fixed, moving and registered volumes, to check if the motion has been properly corrected.

The lesion's position was defined as its center of gravity. As the lesion files have a value of zero everywhere but in the lesion itself, the center of gravity was easily approximated by averaging the position for those pixels with a value different of zero.

The registration was performed with the following algorithms: polynomials from order two to five and the original and improved versions of the demons algorithm, with a $\sigma = 1$ mm gaussian filter after each iteration. Even if it has been found out that the demons algorithms outperform the polynomial ones when it comes to CT co-registration, it is still interesting to confirm that a better CT co-registration leads to a better PET compensation.

4.1.2 Results

The average lesion distance reduction from the moving-fixed image pair to the registered-fixed one is shown in table 4.1.

	Apical (mm)	Base (mm)
Improved demons	2.35 - 2.68 = -0.33	12.61 - 5.84 = 6.78
Original demons	2.35 - 2.43 = -0.08	12.61 - 8.60 = 4.01
Fifth order polynomials	2.35 - 2.49 = -0.14	12.61 - 5.85 = 6.77
Forth order polynomial	2.35 - 2.54 = -0.19	12.61 - 4.60 = 8.01
Third order polynomial	2.35 - 3.18 = -0.82	12.61 - 3.74 = 8.87
Second order polynomial	2.35 - 4.20 = -1.85	12.61 - 10.42 = 2.19

	Central (mm)	Peripheral (mm)
Improved demons	9.07 - 4.04 = 5.03	10.42 - 8.56 = 1.86
Original demons	9.07 - 5.69 = 3.38	10.42 - 8.63 = 1.79
Fifth order polynomials	9.07 - 2.41 = 6.66	10.42 - 5.76 = 4.66
Forth order polynomial	9.07 - 1.89 = 7.17	10.42 - 4.75 = 5.68
Third order polynomial	9.07 - 2.36 = 6.71	10.42 - 3.30 = 7.13
Second order polynomial	9.07 - 7.05 = 2.02	10.42 - 7.95 = 2.47

	Average (mm)
Improved demons	8.61 - 5.28 = 3.33
Original demons	8.61 - 6.34 = 2.27
Fifth order polynomials	8.61 - 4.13 = 4.48
Forth order polynomial	8.61 - 3.45 = 5.16
Third order polynomial	8.61 - 3.14 = 5.47
Second order polynomial	8.61 - 7.40 = 1.21

Table 4.1: Lesion deviation reduction from the fixed-moving image pair to the fixed-registered one, in mm: before - after = improvement.

It can be noticed that the lesions located in the apical region are, as expected, the ones that move the less. In those cases in which the lesion's movement is almost negligible, the registration algorithm does not perform a proper correction. For the other regions (base, central and peripheral), the movements are in general much more prominent (around 10 mm in average).

The averaged values show that the third order polynomial registration seems to offer the best performance, with forth and fifth polynomials just behind. Second order polynomials give poor results, while the demons algorithms are in a disappointing intermediate level. In general terms, 8 mm errors are reduced to 3 mm. This result could be considered as "just satisfactory".

In order to illustrate the results, some images from the lesion that moves the most are shown in figure 4.1. Image a) corresponds to the PET of the lesion in full exhale position. Image b) is the averaged PET all over the breathing cycle. Image c) is an average of the ten volumes after registering them to the exhale position (using third order polynomials). The three lower images are thresholded versions of the upper ones.

The registered version is still blurry and has thus a smaller voxel value (which is related to the standard uptake value, SUV) than the sharp one, but offers a much better better lesion location than the blurry averaged PET in b).

4.1.3 Conclusions

The results show that third order polynomials perform better than the higher order ones and, especially, the demons algorithm. Simpler algorithms, which keep a certain coherence in the deformation, work better than the complex field generated by the demons algorithm.

The reason of why a polynomial of third order outperforms a fifth order one was previously discussed in the "Practical Issues" of the "Image Registration" chapter (the error-order relationship figure 4.2 is repeated here for the sake of comfortability).

The reason for polynomials outperforming the demons algorithm is that the NCAT generated images are very plain (a typical slice is shown in figure 4.3); they do not have many details. They have a lot of voxels with the same value next to each other, which makes the registration software confuse them. It can assume that two pixels are well aligned when they are not (figure 4.4). This can produce a very complex and inaccurate deformation field in the demons algorithm.

It can finally be concluded that the power of the demons algorithm is being wasted, and that polynomials work quite well around the lungs when

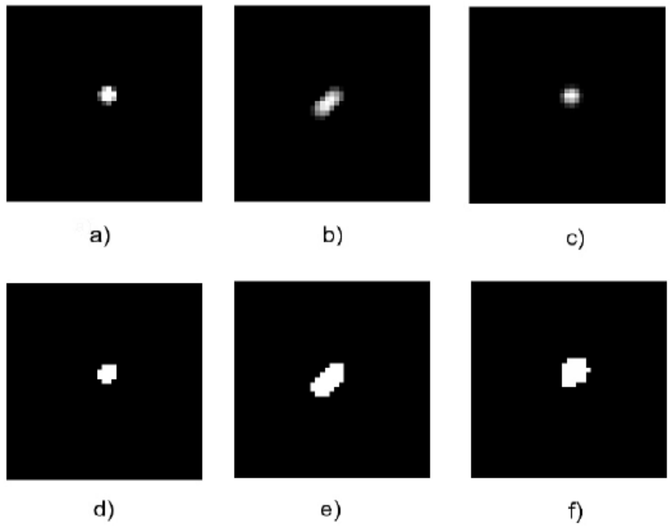

Figure 4.1: Same slice in different PET volumes for the same lesion: a) Static, full exhale b) Average over the breathing cycle c) Average of the registered images over the breathing cycle d) e) and f) are thresholded versions of a), b) and c).

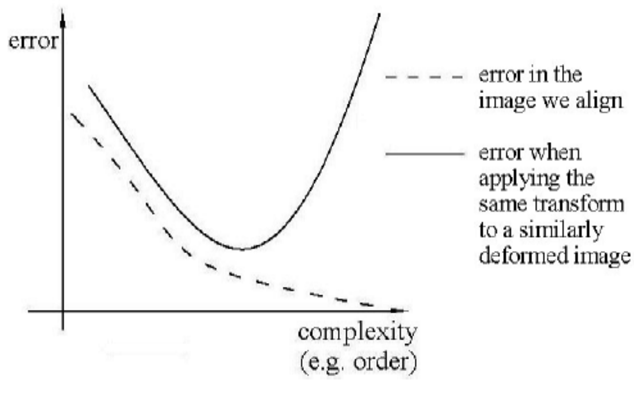

Figure 4.2: Typical error-model complexity curve shape.

Figure 4.3: A typical NCAT-generated transverse slice. The image is quite plain, without many details.

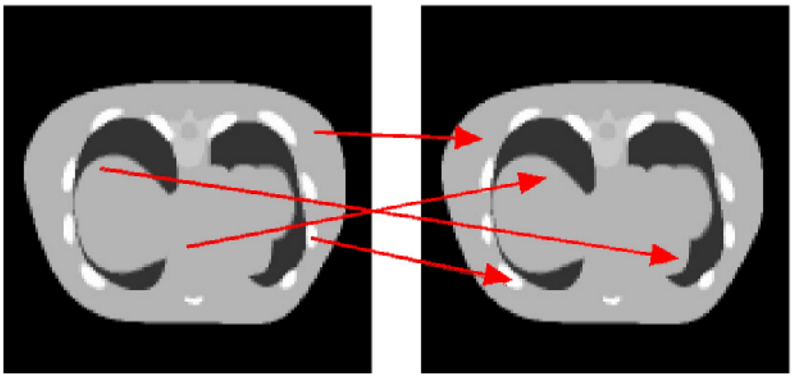

Figure 4.4: Exaggerated, completely erroneous deformation field due to having too many adjacent voxels with the same value.

it comes to phantom images. But one must bear in mind that a phantom is still a fictitious case, and that it does not incorporate much detail. The only way of tackling this problem is working with real images. This is discussed next.

4.2 Compensation with clinical gated PET/CT

In the previous section, the NCAT phantom was used. It has the advantage of being able to generate as many frames as one wants of perfectly matched CT and PET images in the phases one desires. But NCAT also has some disadvantages that have already been discussed above: the generated images are plain and the vector generator tool does not work well. Working with clinical gated PET/CT data is hence the only proper way of testing the compensation algorithm.

4.2.1 Materials and methods

Materials

As the data provided by [36] and used in previous sections to evaluate the registration methods only consists of CT data, [28] were contacted. This group carried out an study whose purpose was "to develop a method for four-dimensional respiration-correlated acquisition of both CT and PET scans and to develop a frame work to fuse these modalities".

A thermometer inserted into the entry of an "oxygen mask" covering the patient's mouth and nose was used as gating device. By selecting slices acquired in the same respiration phase, a 3D reconstruction of the thorax was generated with minimal respiration artifacts. Linear interpolation between two surrounding slices of the raw (350 slices) CT scan was applied when no slice was available at the desired phase.

A sample 4D PET/CT data set (that was the main disadvantage, that it was only one data set) was kindly provided by the authors, and is the base of our ultimate test. In this image set, the CT is gated into 10 different phases and the PET into 16. The correspondence between them is shown in table 4.2. The PET phases were translated to the corresponding CT ones. A corresponding transmission PET (TPET), not gated and thus affected by respiratory motion, is also available.

The PET and CT volumes are not aligned. They do not have the same pixel dimension either. And they do not even represent the same volume: the CT covers the whole thorax, while the PET covers only a certain subregion. The transmission images and the gated PET volumes are in the same coordinates and perfectly matched, though. This fact will makes it possible to match the PET and CT images.

Methods

As explained above, the PET and CT images are unmatched, and it is difficult to register a PET image against a CT. The key fact is that it is much easier to register the transmission image against the CT data than the PET ones, as shapes are much more clear in the TPET (PET images lack anatomical information, figure 4.5). Hence, in order to match the PET and CT data sets, it is possible to register the transmission image to the CT data, store the transformation, and then apply it to the PET data.

CT phase number	PET phase number
3	0
4	1
4	2
5	3
6	4
6	5
7	6
8	7
8	8
9	9
9	10
0	11
1	12
1	13
2	14
2	15

Table 4.2: Correspondence between CT and PET phases. The shift between the data sets is 248 degrees.

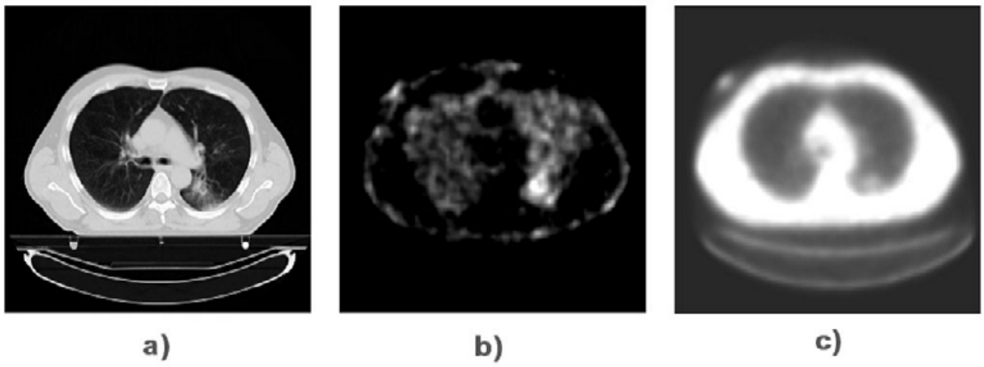

Figure 4.5: Same transverse slice with different modalities: a)gated CT b) gated PET c) transmission image.

The first step was thus to register the transmission TPET image to the CT. As the transmission image includes the respiratory motion (is not gated), an average over the 10 available phases was taken in order to build an equivalent CT volume. After interpolating the transmission image in order to get a volume with the same dimensions and pixel size than the merged CT (zero valued voxels were used to fill the region were there is no transmission image available, as the CT volume was larger), a simple rigid image registration method was utilized to match both volumes.

Mutual information was used as cost function (as different modalities were involved), and only a simple translation transform was optimized; it is assumed that the patient does not move between the CT and PET acquisitions and that the volumes are already in the same scale. Once the registration process was finished, the rescaling and rigid transform process was repeated with the gated PET images in order to express both data sets (PET/CT) in the same coordinate system.

The final step is to co-register the CT images, and evaluate the result of applying the resulting fields to the corresponding PET ones. Let us use the CT volume at phase f as fixed image, and the one at phase m as moving image, and let us call the resulting field $F_{f,m}$. If $F_{f,m}$ is now applied to the PET image at phase m, one would expect a PET image at phase f.

The registration was performed with the same algorithms as in the previous section (NCAT). The polynomials (apart from the demons algorithms, of course) will still be investigated in order to try to confirm that a better CT co-registration leads to a better PET compensation (it is important to bear in mind that one may fall into "overfitting").

The similarity between the fixed and registered images and the lesion's center of gravity location improvement are used as measures of how successful the method has been. The calculation of the lesion's center of gravity was performed in a simplified manner, by averaging the coordinates of the pixels above a threshold (which was arbitrarily defined as the 33% of the image's maximum value) .

4.2.2 Results

To illustrate the similarity range of the gated PET images and have a pattern to compare later results, the similarity measure between PET at phase 0 and the other phases is illustrated in figure 4.6 (before any correction being applied). A cyclic behavior, direct consequence of the breathing's periodicity, can be observed. One can infer that a SSD of less than 3000 or an NCC of 0.92 can be interpreted as "very similar images" in this data set.

Figure 4.7 shows the average similarity measures and lesion deviations

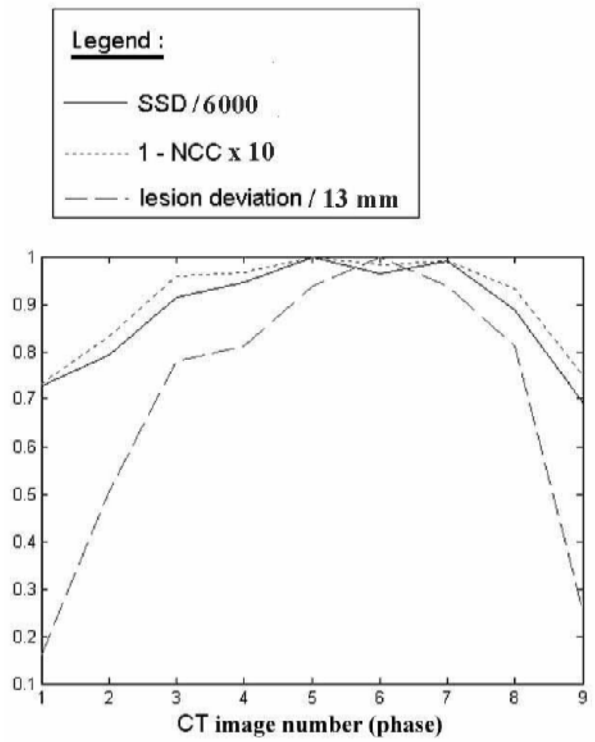

Figure 4.6: Normalized similarity measure values between PET in phase 0 and the other ones: a) SSD, b) 1- NCC, c) lesion deviation.

between the PET phase 0 and the compensated PET images from the other nine phases (that is, the result of transforming the other nine phases with the transforms obtained from the CT data).

If ones pays attention to the SSD similarities first, it can be noticed that a third order polynomial gives the best results. The “too complex transform” effect can be observed again with the higher level polynomials, while the demons algorithms are at an intermediate level. The third order polynomials are the only one that get clearly under the SSD = 3000 level, which was defined as “very similar”.

On the other hand, if the NCC similarity measure is considered, the demons algorithms outperform the polynomials. Among the latter ones, the same “overfitting”effect as before can be observed with the fifth order ones. The original demons algorithm performs better that the improved one, but none of them reaches the “very similar” level of NCC = 0.92 (just 0.90).

The most interesting result is presented in figure 4.7-c. The polynomial algorithms are not capable of correcting the lesion position that much, while the demons algorithms, and especially the improved one, are able to do it (which is consistent with the NCC similarities and the results from the previous chapter, although inconsistent with the SSD similarities). An (average) initial deviation of 12 mm is reduced to 7.5 mm by the original method and to 5.7 mm by the improved one. Sample images from this correction process are shown in figure 4.8.

4.2.3 Conclusions

The results show that the demons algorithms in general, and the improved one in special, can partially compensate for the lesion deviation due to the patient’s respiratory movements.

Checking a region of some registered PET images with different methods, and change the contrast to show darker details, it can be noticed that the improved demons algorithm is prone to “scratch” the image (figure 4.9). This may be the reason of the poor SSD similarity (even if the NCC similarity and the position compensation are satisfactory).

One can also notice that demons algorithm is also prone to deform the lesion. There are two reasons for this. The first one is that the PET and CT images are not perfectly aligned. In that case, a deformation vector that is not exactly the correct one is being applied to each point. This effect is less prominent if polynomials are used, as the vectors change more smoothly.

The second reason is that the smoothing filter variance, which was set to $\sigma^2 = 1mm^2$ (following the proposal in [25]), is not high. That would let the

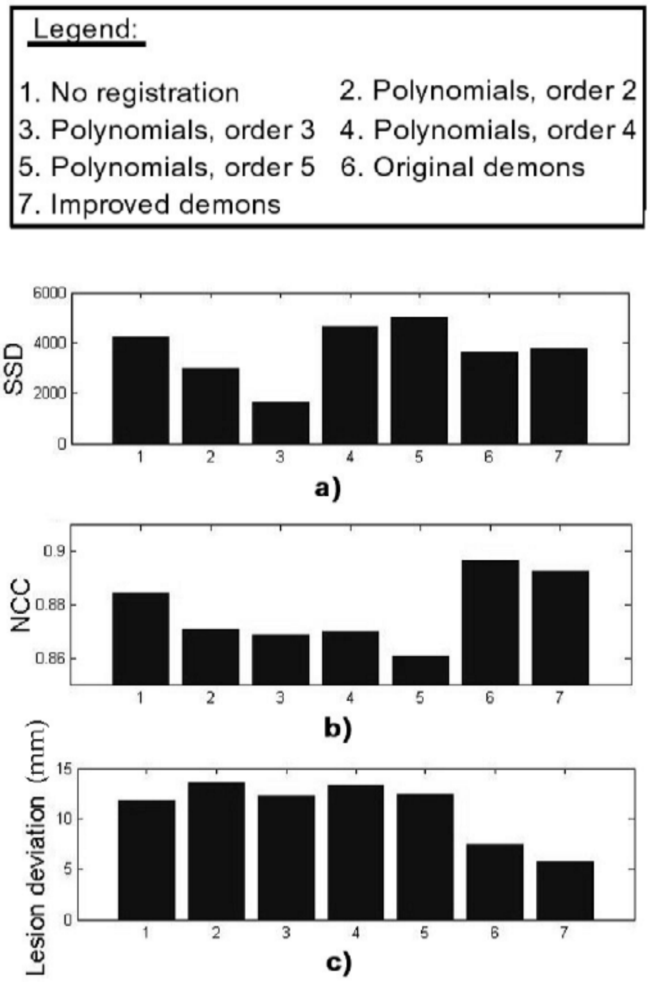

Figure 4.7: Averages a) SSD, b) NCC, and c) lesion deviation for the PET phase 0 and compensated PET images, with seven different registration methods.

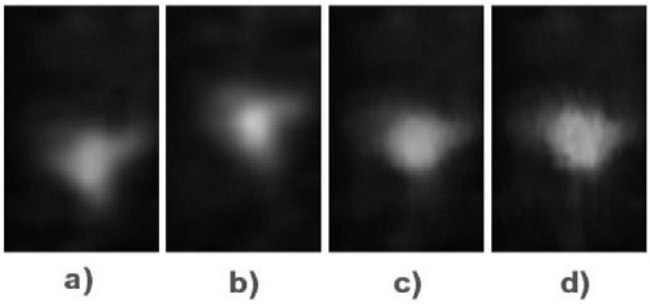

Figure 4.8: PET image compensation process: a) lesion in phase 0 b) lesion in phase 7 c) compensated with original demons algorithm d) compensated with improved demons algorithm.

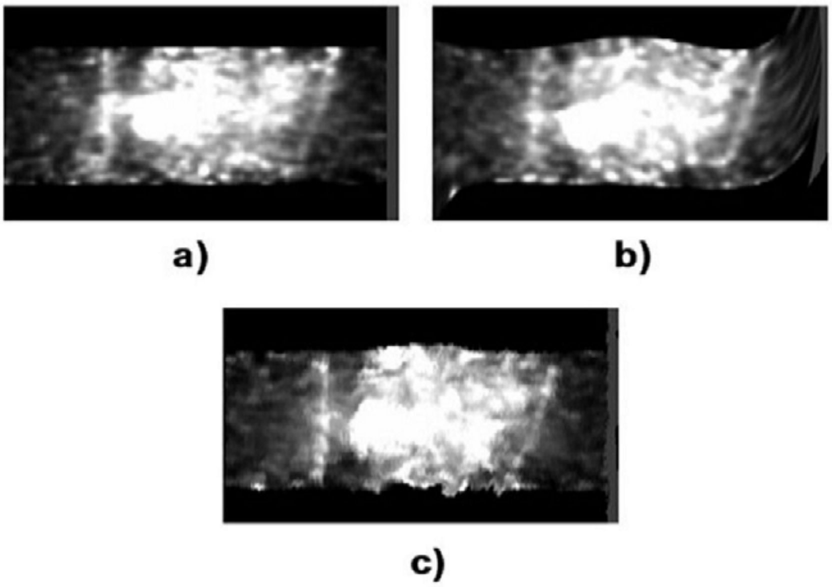

Figure 4.9: a) Fixed image b) registered with a third order polynomial c) registered with the improved demons algorithms. Please notice the scratched appearance of the image.

field vary rapidly, which explains explain the “scratched” appearance of the registered image. A sample field is shown in figure 4.10.

In order to finish, it can be concluded that, when it comes to clinical data and evaluating the lesion position correction, the demons algorithm (especially the improved version) outperforms the polynomial methods. This was expected, as this was also the case when co-registering CT images in different respiratory phases in the previous chapter (whose results become consistent with this chapter's). This is the reason why the focus will be set on this algorithm from now on, aiming at improving the alignment stage and optimizing the smoothing filter variance.

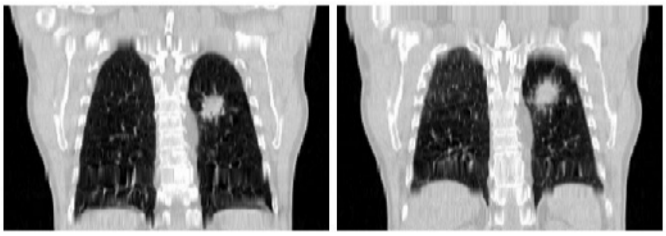

a)

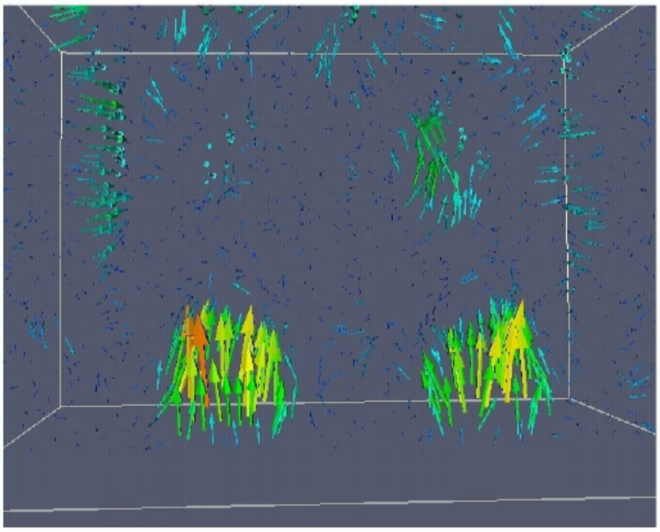

b)

Figure 4.10: a) Coronal view from CT images in different respiratory phases b) Deformation field sampled along the same plane as the CT image, calculated with the improved demons algorithm. Mind the big deformation in the lower lobes and around the lesion, as well as the messy distribution of small vectors around the rest of the image.