Chapter 5

Optimizing the demons algorithm

Once the improved demons algorithm has been chosen as the most appropriate one for registering the images in this thesis, its application on this work will be optimized

The first step will be to improve the alignment stage between the transmission PET image and the averaged CT one. Afterwards, the smoothing filter variance will also be optimized. One can finally try modify the number of iterations at each resolution level in the pyramidal scheme. All the work is based on the clinical data set from [28].

5.1 CT-TPET alignment

If one superimposes a certain CT image and the corresponding aligned PET one in the same phase (aligning them with the same algorithm as in last chapter), it is noticed that the alignment is poor. A graphical example is show in figure 5.1. There is clearly a lesion that is located in the CT image in a higher position than in the PET one.

A new algorithm, based on second order polynomials, is proposed in this section. The steps are:

1. Preprocessing the blurry CT image in order to make it as similar as possible to the TPET one. This can be accomplished by applying a wide gaussian blurring kernel on it. A Gaussian filter with a spatial variance of $\sigma^2 = 60\ mm^2$ was used. A sagittal slice from the filtered blurry CT, along with the corresponding slice from the TPET image, is represented in figure 5.2.

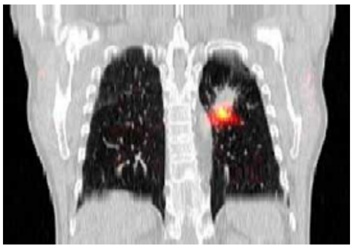

Figure 5.1: CT image (shown in greyscale) and corresponding aligned PET image (in red) superimposed to each other, using last chapter's alignment algorithm.

2. A second order polynomial registration was performed with the corresponding AIR program. The standard deviation of ratio images cost function was used. As described in the program instructions, this cost function has the advantage of being independent of image intensity, so that image intensities can be poorly matched and the registration will not be adversely affected. This cost function is suitable for intermodality registration, which is now the case (even if it has been tried that the images become as similar as possible).

3. The calculated transform is applied on the PET images.

If a CT image and the corresponding PET in the same phase (after aligning it with the new method) are now superimposed, the visual impression is much better (figure 5.3). If the deformation fields calculated in the previous chapter are now applied, the similarity measures between the PET images improve, as it can be appreciated in figure 5.4. Even though the SSD measure is worse (strange results had also been obtained with this metric in the previous chapter), the NCC and, what it most important, the tumor's center of gravity deviation, are significantly better.

Even if the alignment has improved, it is still worth to mention that the results would be better if the images were acquired by an hybrid PET-CT scanner. In the available data set, the CT and PET images were acquired in different scanners and in different time instants. That is the reason why a simple resizing / translation algorithm is not enough for aligning the images. With an hybrid CT-PET scanner, this alignment step could be completely omitted; the term "hardware matching" is usually used to describe the technique. These results have thus a big margin for improvement.

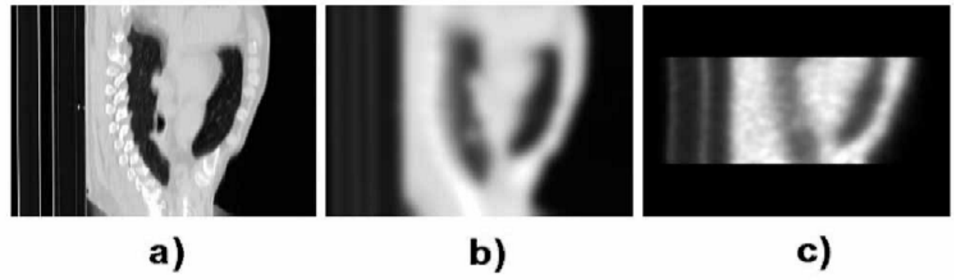

Figure 5.2: a)Time-averaged blurry CT image b) Very blurred version of the previous one c) TPET image to register against b)

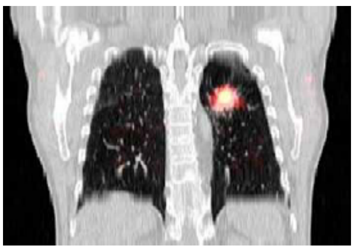

Figure 5.3: CT image (shown in greyscale) and corresponding aligned PET image (in red) superimposed to each other, using the new aligning algorithm. The results are much better than with last chapter's algorithm.

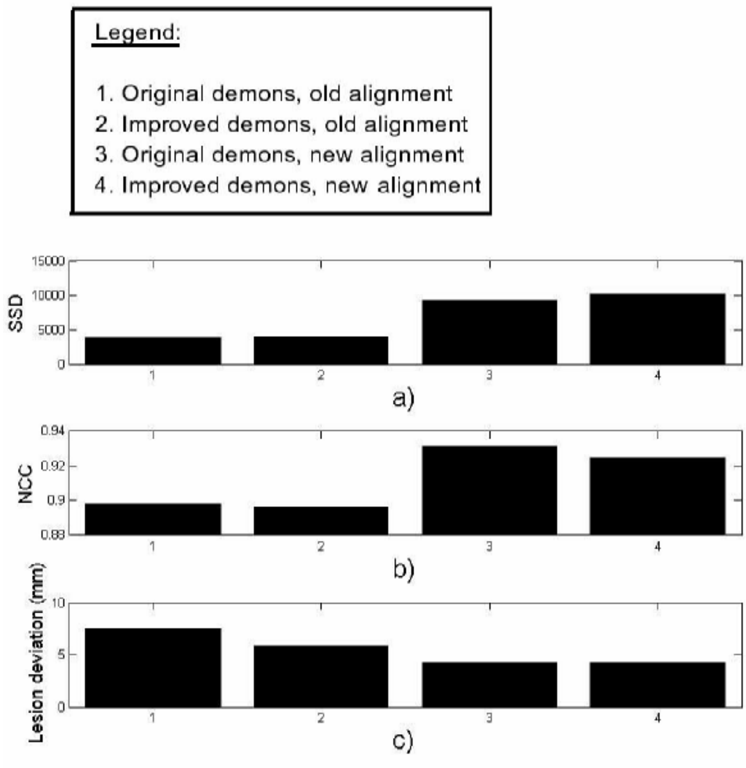

Figure 5.4: a) SSD, b) NCC, and c) lesion deviation for the fixed and registered images, both with the old and the new alignment method.

5.2 Optimizing sigma with time constraints

Another of the filter parameters to be optimized is the variance of the Gaussian smoothing filter applied after every iteration in the demons algorithm. This filter, along with with the pyramidal processing, are the utilized mechanisms in order to try to make the resulting deformation field look like a real movement field. Applying this filter prevents the algorithm from generating very different movement vector in points that are close to each other, which would represent an unreasonable deformation. Using a too high variance takes away flexibility from the transform, while a too low one may generate an unrealistic field or even prevent the algorithm from converging.

An important consequence of modifying this value is that the converge speed can be affected. The higher the spatial variance, the less flexible the transform and the less iterations are necessary to converge. So, if one wants to be fair with all the candidates, one must use the same number of iterations for all of them in order not to give any time advantages. But if the focus is to be set on quality, without any time constraints, the possibility of letting the algorithm iterate for a longer time can be considered.

It is also important to mention that the filtering time after each iteration is neglected, that is, it is assumed that this time is much smaller that the time it takes to deform the images and calculate the gradient in each point. in this case it can be assumed that two registration processes with the same number of iterations at each resolution level take the same time. This is only true if they use the same variance value for the smoothing filter; when a higher variance is used, our program requires a bigger kernel to successfully approximate the Gaussian filter and the calculations take a longer time. But if it is assumed that this time can be neglected when adding it to the gradient calculation and image deformation times, as mentioned before, it is possible to assume the registration times to be approximately equal.

Maintaining a 14-10-6 scheme (iterations in each resolution level) for the original demons algorithm and a 7-5-3 for the improved one (as in the previous chapter), which is enough for a value of σ of 1-2 mm, and registering the full exhale phase to the full inhale one, the average similarity measures shown in figure 5.5 are obtained. It can be noticed that the improved demons algorithm peaks at approximately 2 mm for all of the quality measures, while the original algorithm does it around 1 mm (although it shows a strange behavior in the lesion deviation).

In order to improve the resolution around the area where the curves reach their minimums, the experiments were repeated with smaller increments around the mentioned values of one and two millimeters. The new results are shown in figures 5.6 and 5.7. The improved algorithm peaks

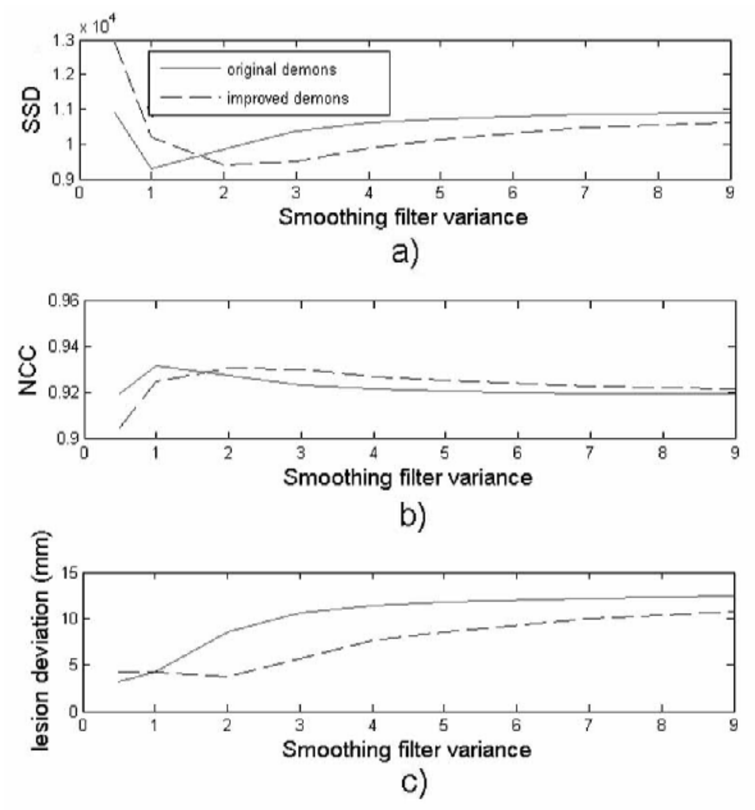

Figure 5.5: a) SSD, b) NCC, and c) lesion deviation depending on the smoothing filter variance for both the original and the improved demons algorithm. The full exhale phase was registered to the full inhale one.

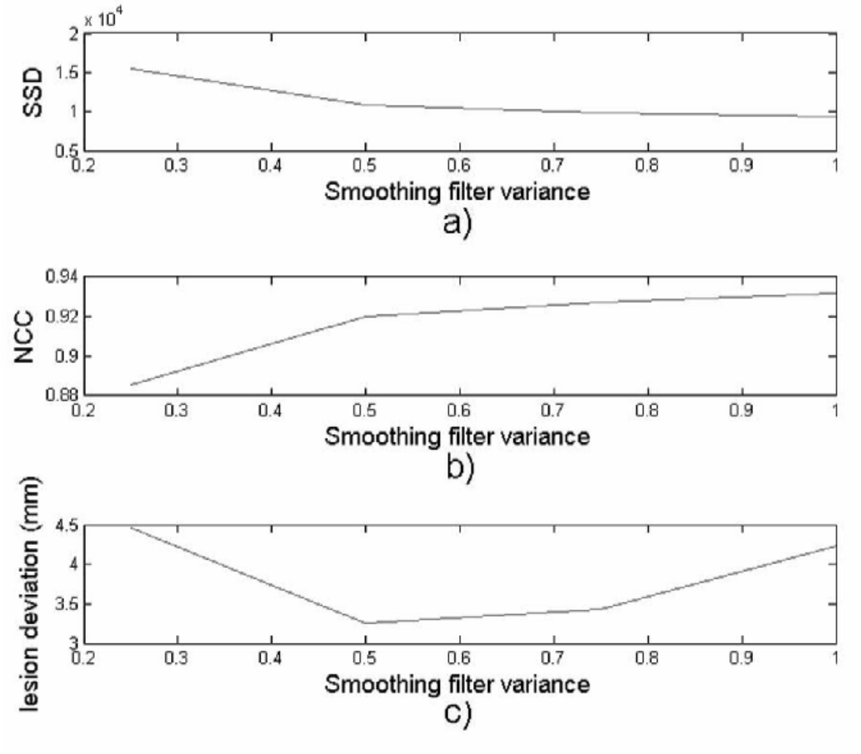

Figure 5.6: a) SSD, b) NCC, and c) lesion deviation depending on the smoothing filter variance for the original demons algorithm.

clearly at approximately 1.75 mm, while the original one does it at approximately 1 mm for the similarity measures and around 0.5 mm for the lesion deviation. It can be noticed that, in the registered image (figure 5.8), $\sigma = 0.5$ mm really deforms the tumor a lot. Hence, a compromise of 0.75 mm is a good solution.

After optimizing the filter variance, the mean tumor deviation has reduced to 3.5 mm. Considering that the original value was 13 mm, the image misalignment (as an hybrid PET/CT scanner was not available), and that the pixel size in the z direction is of 3 mm (1.76 mm in the $x-y$ plane), this is a very promising result. The following step is to try to further improve the results by eliminating the processing time constraints (although it may not be practical). This is done in the following section.

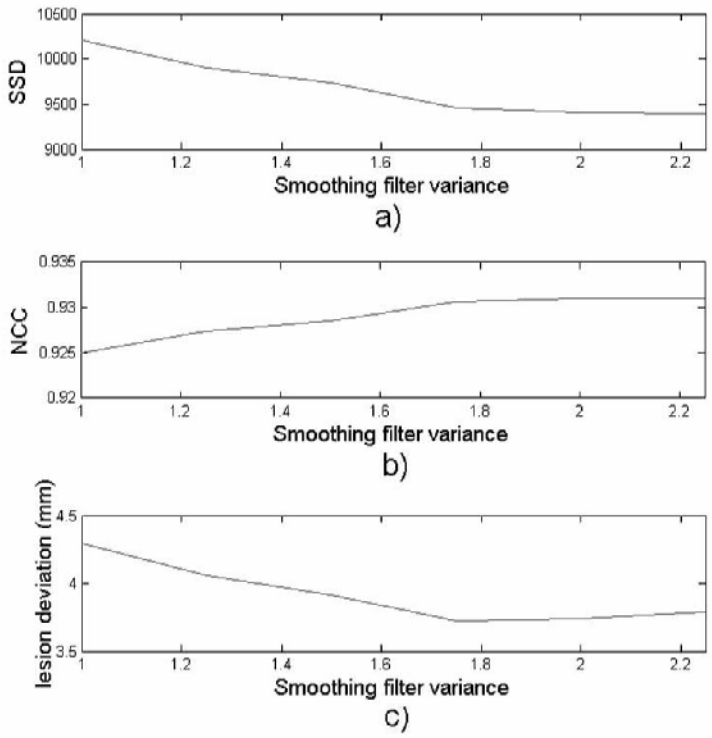

Figure 5.7: a) SSD, b) NCC, and c) lesion deviation depending on the smoothing filter variance for the improved demons algorithm.

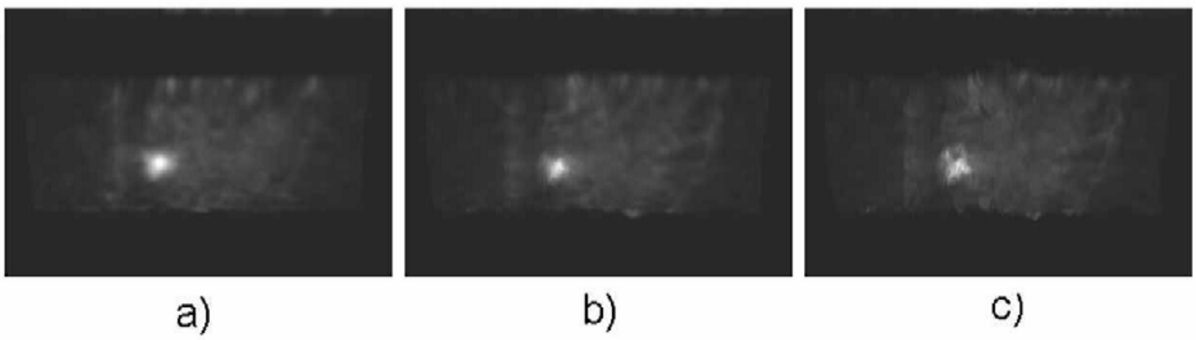

Figure 5.8: a) Original PET image in phase 0 b) Compensated image from phase 5 and variance of 1 mm c) Variance of 0.5 mm.

5.3 Optimizing sigma without time constraints

In the previous section, a fixed number of iterations was used in the demons algorithm. The decision was based on the observed metric's behavior for several images, using the value $\sigma = 1$mm proposed in [25]. Now, different values of σ, that may require more or less iterations to converge, are inspected. Some of them would require more time, but it is still interesting to find out if the quality of the results can increase. Even if the registration time becomes too long, one must consider that computers evolve really quickly (something that is computationally expensive today may be affordable in the future) and that pyramidal processing provides with preliminary approximate solutions in a small time period.

5.3.1 Original demons algorithm

Observing the metric evolution for the original algorithm in figure 5.9, it could be possible to improve the SSD similarity results by taking some more iterations (especially in the highest resolution, which is unfortunately the most computationally expensive one). It can be observed that a higher value for σ leads to a shorter convergence time, too. This was expected, as the higher the variance, the less flexible the transform, and the less iterations to converge.

The strange peaks in the figure when the variance is low deserve also being remarked. They coincide with the points in which the resolution level is changed. When a field is interpolated to the next level, more mistakes are made than in the previous one. This is due to “overfitting” (an already commented effect).

In order to try to find an optimum value of sigma, a variant of the registration program was designed. In that variant, instead of using a fixed number of iterations at each level, the processing is finished after N_i iterations without improvement at level i, where Ni is an user-defined variable. The Ni values were arbitrarily set up to three, two and one iterations for the lowest, intermediate and highest resolution level, respectively.

The SSD similarity metric against the iteration number is plotted for two sample image pairs in figure 5.10. One can notice that the number of iterations becomes significantly larger than the original fixed number of iterations, and that there are tens of them in the highest resolution level. This makes the registration time much longer and, apart from that, is probably leading to“overfitting”. That is the reason why the original number of iterations can be considered to be balanced.

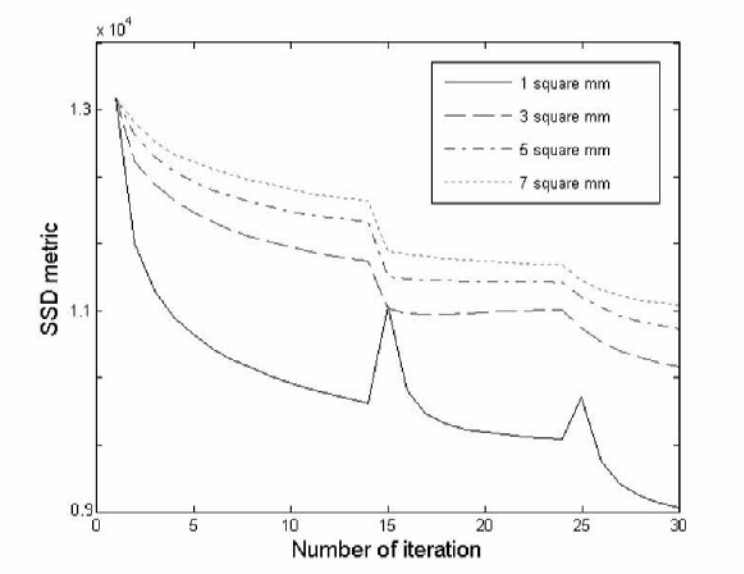

Figure 5.9: Average SSD metric for different smoothing filter variances against iteration number (original demons algorithm).

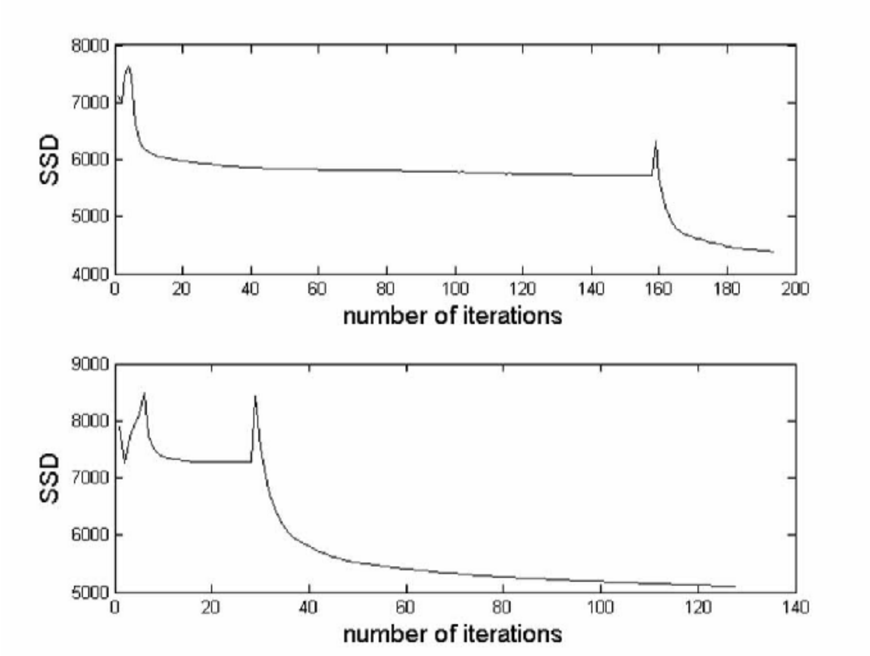

Figure 5.10: Long registration with the original demons algorithm. The iteration process is stopped after 3, 2 and 1 iterations without improvement in each resolution level.

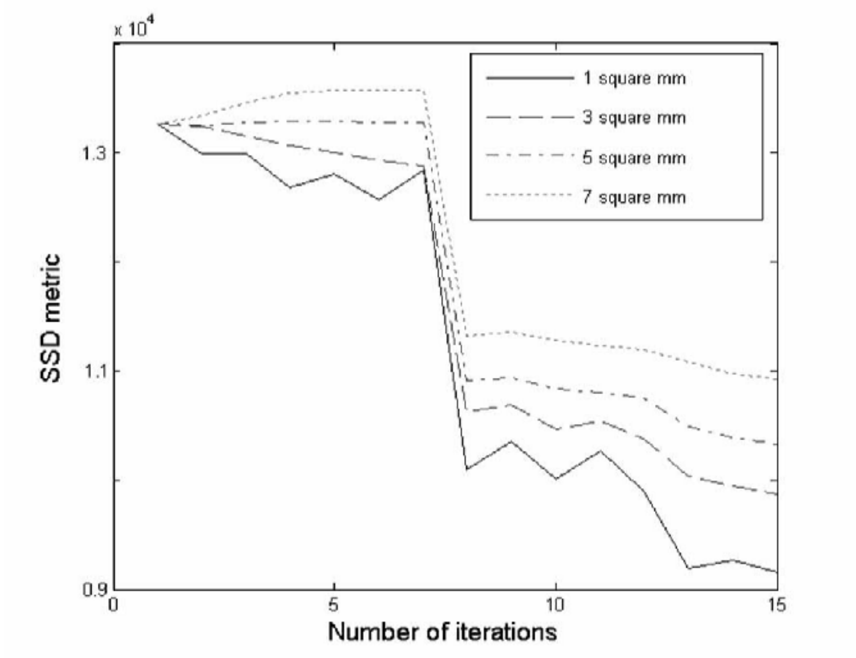

Figure 5.11: Average SSD metric for different smoothing filter variances against iteration number (improved demons algorithm).

5.3.2 Improved demons algorithm

The conclusions that can be extracted from the SSD metric evolution for the improved algorithm (figure 5.11) are different to the ones for the original algorithm. It is still true that higher variances lead to faster convergence, but the behavior is much more irregular. It can be noticed that, if the variance is high, the iterations in the lowest resolution level actually worsen the SSD similarity. Another important difference is that the resolution changes suppose a large improvement in the metric. This algorithm leads to a field that does not "overfit" as in the previous case, being in principle a better approximation of the real deformation field.

It can finally be commented that the results could be improved by taking some more iterations (especially in the highest resolution, as before, but now it is even more computationally expensive). And this time one does not need to be careful with the "overfitting" effect as before. In figure 5.12, the SSD metric against the iteration number is shown for a sample pair of images. 7, 5, and 15 iterations were performed, that is, the usual fixed amount plus 12 extra iterations in the highest resolution level.

The final metric value can definitely be improved by taking this extra

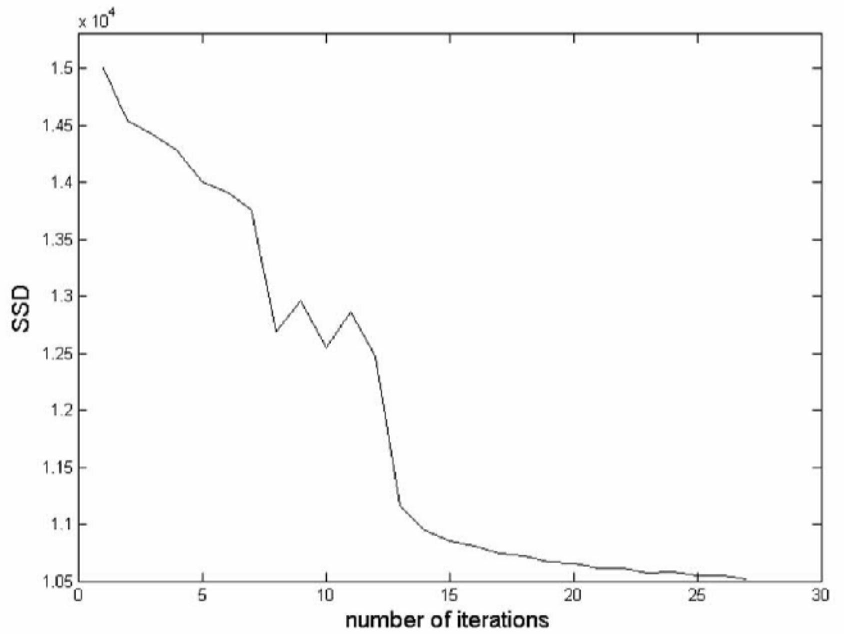

Figure 5.12: Long registration with the improved demons algorithm. Twelve more steps that usual are taken in the highest resolution level.

	7-5-15	7-5-3
SSD	9631	10171
NCC	0.9289	0.9246
lesion deviation (mm)	3.45	3.86

Table 5.1: Result comparison for a 7-5-3 and a 7-5-15 iteration scheme for the same case as the figure 5.12.

iterations. The discussion of that if this is worth depends on the application speed requirements and on the speed of the machine where the algorithms are run. The similarity metrics for the 7-5-3 and 7-5-15 schemes are shown in table 5.1. There is a fair improvement, but the processing time increment is also very large. In this thesis the 7-5-3 scheme was chosen, as many different experiments were to be performed and the registration speed was preferred to be as high as possible.

5.3.3 Conclusions

It can be concluded that one should focus on optimizing the number of iterations for the optimal variance value (or a slightly smaller one), rather than on finding out if using other variance values can produce higher similarities with a different number of iterations. If a larger amount of iterations gives worse results for a higher variance, which is supposed to converge faster, that setup should not be considered. That is the reason why the values of $\sigma = 0.75$ mm (original algorithm) and $\sigma = 1.75$ mm (improved algorithm) are chosen.

5.4 Conclusions: final method

A new TPET - averaged CT registration method was proposed in this chapter. Along with it, using a gaussian spatial variance of 0.75 mm and a 14-10-6 iteration scheme for the original demons algorithm (1.75 mm and 7-5-3 for the improved one) is the base of our ultimate compensation method. Its performance will be investigated more deeply in the next chapter.