

Guided Filtering and Partial Volume Correction for Positron Emission Tomography

Dóra Varnyú (varnyu.dora@gmail.com), László Szirmay-Kalos (szirmay@iit.bme.hu)

Department of Control Engineering and Information Technology, Budapest University of Technology and Economics, Budapest, Hungary

Abstract

Positron emission tomography (PET) is a nuclear medicine imaging technique that is used to observe tissue metabolism by reconstructing the spatial distribution of the injected radioactive tracer. Due to constraints on the time and the radiation dose of the examination as well as limited scanner sensitivity, PET images usually suffer from a high level of noise. This paper focuses on the application of the guided filter for PET image denoising. After proposing several different guidance images, guided filter variants are compared with the median, the Gaussian and the bilateral filter in terms of image quality and speed. For dynamic PET reconstructions, a new approach, the parametric filtering is conceived, in which filtering is performed on the parameters of the kinetic model describing the radiotracer concentration. Finally, an efficient, guided-filter-based partial volume correction (PVC) method is proposed to restore accurate activity values that are blurred due to the partial volume effect.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Filtering

1. Introduction

Positron emission tomography (PET) reconstructs the spatial distribution of the radiotracer injected into the patient. The tracer gets absorbed by tissues in proportion to their metabolic rate, therefore regions with fast metabolism such as the brain, the heart or malignant tumors appear bright in the image.

During reconstruction, the measured volume is divided into a 3D grid of homogeneous voxels. The task is to determine the amount of radiotracer in each of the voxels based on the emission detected by the tomograph. This is accomplished by the maximum-likelihood expectation-maximization (ML-EM) iterative procedure [13], which alternates forward projections and back projections to find the voxel activities that have the highest probability to have generated the measured data.

There are two types of PET reconstruction, static and dynamic. With a static reconstruction, the goal is to visualize a single snapshot of the patient's body, whereas dynamic reconstructions focus on monitoring biochemical changes taking place in the body, such as the progression of a disease or the absorption of a certain drug in the organs.

In dynamic PET, the space-time dependence of the tracer concentration is reconstructed. The measurement time is divided into independent time frames. Radiotracer concentration is described by a kinetic model of a predefined algebraic form that assigns a parameter vector to each voxel. The goal of the reconstruction is to calculate these parameter vectors [16,19]. In this paper, the two-tissue compartment model [20] and the spectral model [18] are examined.

1.1. Previous work

PET images generally suffer from high noise due to constraints on the time and the radiation dose of the examination as well as limited scanner sensitivity, which can result in misdiagnosis. Noise can be suppressed by introducing a regularization term into the target function of the ML-EM optimization [10], but defining the right parameters is challenging because their optimal values depend greatly on the measured data and the reconstruction settings. Also, if the regularization term is not convex, optimizing the modified target function becomes difficult [22]. Alternatively, noise can be reduced by early stopping [5] or by filtering [6, 10]. In case of early stopping, the signal has not yet converged to

the optimum, so the image may not be detailed enough. In clinical application, filtering has become more common.

Filtering can be performed at the end of the reconstruction, or after each iteration by incorporating it into the reconstruction scheme. This latter technique is also known as the method of sieves [9, 14] and its main aim is to keep the result of the ML-EM optimization within the set of acceptable solutions.

The most commonly used filter is the Gaussian filter due to its simplicity and efficient implementation. However, the Gaussian filter smooths not only the noise, but also the edges, which can lead to the loss of important clinical information. Edge-preserving filters like the bilateral filter [11, 17] may provide a better alternative. However, the bilateral filter was reported to suffer from the gradient reversal problem, the introduction of false edges in certain scenarios [4].

Previous works have applied the guided filter, a fast, edge-preserving operator developed by He, Sun and Tang [4] to suppress noise in dynamic PET reconstructions [7, 21].

Guided filter produces its output based on the contents of a guidance image. This makes it possible to use additional anatomical information and previous reconstruction results to improve the contrast of the reconstructed image. Furthermore, the guided filter is also devoid of the gradient reversal problem observed with bilateral filtering.

Lu et al. [7] proposed a guidance obtained by summing the activity images of frames, while Yan et al. [21] used an MRI image to guide the filtering. Both their works focused on dynamic reconstructions and examined only one guidance.

1.2. Proposed solution

In this paper, the effect of guided filtering is examined in both static and dynamic reconstructions. New guidances are proposed for both types of reconstruction, which are constructed either by pre-processing the input image or by using multiple available data together in a multi-channel guidance. The proposed guided filter variants as well as the variants from previous works are compared with the median, the Gaussian and the bilateral filter in terms of image quality and speed. The difference between post-reconstruction filtering and the method of sieves is also examined. Furthermore, for dynamic PET reconstructions, a new approach called parametric filtering is elaborated, in which filtering is performed on kinetic parameters instead of on activity images as before.

In addition to high noise, another challenge in PET imaging is the partial volume effect (PVE), the phenomenon of a tissue's activity affecting the reconstructed activity concentrations of the surrounding tissues, also called spillover of activity. The main reason for this is the limited resolution of the imaging system and the heterogeneity of the voxels. To restore accurate concentrations, the partial volume

effect should be incorporated into the reconstruction model, and correction should be applied. In the paper, an efficient, guided-filter-based correction method is proposed and analyzed.

Our implementation is accelerated by parallel execution on the graphics card using the CUDA platform [1, 16].

2. The Guided Filter

The key assumption of the guided filter is a local linear model between a guidance image G and the output image Q [4]. That is, the output image Q is a linear transform of G in a window ω_k centered at voxel k :

$$Q_i = a_k G_i + b_k, \quad \forall i \in \omega_k, \quad (1)$$

where a_k and b_k are linear coefficients presumed to be constant in the local window ω_k .

Guided filtering aims to find the output image that is the closest to the input, while maintaining the local linear model in (1). To achieve this, the following cost function is minimized:

$$E(a_k, b_k) = \sum_{i \in \omega_k} (\|Q - P\| + \epsilon a_k^2) \quad (2)$$

$$= \sum_{i \in \omega_k} ((a_k G_i + b_k - P_i)^2 + \epsilon a_k^2), \quad (3)$$

where ϵ is a regularization parameter preventing too large a_k values. The solution is given by:

$$a_k = \frac{\frac{1}{|\omega|} \sum_{i \in \omega_k} G_i P_i - \mu_k \bar{P}_k}{\sigma_k^2 + \epsilon}, \quad (4)$$

$$b_k = \bar{P}_k - a_k \mu_k, \quad (5)$$

where μ_k is the mean and σ_k^2 is the variance of G in ω_k , $|\omega|$ is the number of voxels in ω_k and $\bar{P}_k$ is the mean of P in ω_k . Thus the output is calculated as follows:

$$Q_i = \frac{1}{|\omega|} \sum_{k|i \in \omega_k} (a_k G_i + b_k) = \bar{a}_i G_i + \bar{b}_i, \quad (6)$$

where $\bar{a}_i = \frac{1}{|\omega|} \sum_{k \in \omega_i} a_k$ and $\bar{b}_i = \frac{1}{|\omega|} \sum_{k \in \omega_i} b_k$ are the average coefficients of all local windows covering voxel i .

The key question is what to choose as a guidance to keep tissue outlines as sharp as possible while suppressing noise. Applicable guidances depend on the type of the reconstruction, that is, whether static or dynamic reconstruction is performed.

3. Methods and Materials

3.1. Noise Reduction in Static Reconstructions

In the simplest case, the input itself is the guidance (*Input Guided Filter, IGF*). In many cases, this can produce

good results. However, if the input is very noisy, the high-frequency components will appear on the output. To avoid this, we propose preprocessing the input image in two steps before using it as a guidance. First, high-frequency components are eliminated by Gaussian filtering, then edges are enhanced by high-boost filtering (*Enhanced Input Guided Filter, EIGF*).

One of the major benefits of guided filtering is that it allows the use of anatomical information, since the guidance can be the output of another imaging modality, such as CT (computed tomography) or MRI (magnetic resonance imaging), as Yan et al. proposed [21]. Using an anatomical image, the filtering can take into account tissue boundaries (*Anatomically Guided Filter, AGF*). However, it is not always appropriate to use the anatomical image directly as a guidance, since the images of the different modalities carry different information.

To account for this difference, we propose a guidance that uses the other modality image only indirectly, as the guidance of guided bilateral filtering [2] on the input image (*Anatomically-Guided-Input Guided Filter, AGIGF*):

$$G_k = \frac{\sum_{i \in \omega_k} P_i \cdot g(\|i - k\|) \cdot g(A_i - A_k)}{\sum_{i \in \omega_k} g(\|i - k\|) \cdot g(A_i - A_k)}, \quad (7)$$

where P is the input image, A is the anatomical image, G is the proposed guidance, k is the voxel being filtered, ω_k is the window centered at voxel k , g is the Gaussian function, and $\|i - k\|$ is the Euclidean distance of voxels i and k . The resulting image carries information on anatomical tissue boundaries and is also correct in terms of intensity ratios.

Guided filter can also be extended to multi-channel guidances [4]. For static PET reconstructions, we propose a dual-channel guidance, where the first channel is the enhanced input image and the second is the anatomically guided input (*Dual-Channel Guided Filter, DCGF*). Using the two channels together provides such additional information for the filtering that would not be available using the guidances separately.

3.2. Noise Reduction in Dynamic Reconstructions

In case of dynamic reconstructions, parameter vectors of a predefined kinetic model are determined for each voxel. After executing forward and back projection for each time frame, parameters are fitted to the activity values obtained.

Conventionally, filtering is performed on the activity images of the different frames. In this paper, we propose another approach in which filtering is performed only after parameter fitting, on the voxel arrays containing the resulting kinetic parameters. Hereinafter, this latter approach will be referred to as *parametric filtering* and the former as *activity filtering*.

Whichever solution we choose, the guidances described

for static reconstructions can be applied for dynamic reconstructions as well. However, in dynamic reconstructions, it can be exploited that the same object is measured in each frame. Assuming the object does not move during the scan, tissue boundaries should be in the same location in each frame.

Lu et al. [7] proposed a guidance obtained by summing the activity images of all frames. Note that the resulting composite image is essentially what the output of the static reconstruction would have been, that is, if we did not divide the measurement time into frames. Therefore, we call this variant the *Static Reconstruction Guided Filter (SRGF)*.

However, if frames are summed, the information on the time of the activity's detection is lost. It is possible that one of two adjacent regions had high activity in the first half of the measurement and the other in the second half. In this case, summing activities over the entire measurement time can cause the boundary between the two regions to disappear, making the regions blend together. To tackle this problem, we propose grouping the frames and creating a multi-channel guidance where each channel is the sum of the activity images of a group. It would also be possible to use each frame image as a channel itself, without grouping. However, multi-channel guided filtering requires inverting a matrix of the size equal to the square of the number of channels. Since matrix inversion is an expensive operation, it is worth keeping the number of channels as low as possible. For this reason, we used three channels (*Triple-Channel Guided Filter, TCGF*).

3.3. Partial Volume Correction

Partial volume effect (PVE) manifests as a blurring operation on the reconstructed image. Mathematically, it can be described by a convolution with the point spread function (PSF), the system's response to a point source of activity [21].

To compensate for PVE, we modify the aim of the guided filtering to minimize the difference between the output image convolved with the PSF and the input image, that is, $\|(PSF \otimes Q) - P\|$ instead of $\|Q - P\|$, where $\otimes$ denotes the convolution operator. This can be achieved by convolving the guidance with the PSF before calculating coefficients a_k and b_k as described in (4) and (5), then using the original guidance to produce the output $Q_i = \bar{a}_i G_i + \bar{b}_i$. The inverse effect of the blurring, or in other words, the sharpening will appear on the resulting image.

4. Results

4.1. Noise Reduction in Static Reconstructions

Numerical evaluation of the filters was performed on the NEMA NU 4-2008 preclinical phantom [3] (Fig. 1).




Figure 1: Axial slices of the NEMA NU 4-2008 phantom

4.1.1. Image quality

On the NEMA phantom, image quality is characterized by how well the individual rods, in particular the thinnest rod, are reconstructed. This is described by the recovery coefficient (RC), which is the quotient of the reconstructed and the actual activity concentration in a rod.

A filter is considered good if the RC does not decrease significantly with increasing blur strength. Blur strength is calculated as the standard deviation of the activity values in the uniform region. Standard deviation can also be given as a percentage, if it is divided by the average activity and then multiplied by 100.


Fig. 2 shows how the RC values of the examined filters change as a function of blur strength. Different blur strengths were achieved by changing the parameters of the filters, such as the radius of the filter window.

It can be seen that Gaussian filtering significantly decreased the visibility of the rod, but the median filter yielded even worse RC values. With the edge-preserving bilateral filter, the RC decreased to a lesser degree.


Guided filtering achieved extremely good visibility. Of its various types, the input guided produced the lowest RC values, but even these were significantly higher than that of the median, the Gaussian and the bilateral filters.

With this homogeneous phantom, there was no significant difference between using the anatomical image directly as a guidance or only indirectly, in a guided bilateral filtering. On average, the RC values obtained in the latter case were 0.002 higher than in the former, with the same parameterization.

Using the enhanced version of the input image as a guidance yielded very good results. When the blurring was slight, not only did it not reduce, but it actually improved the visibility of the rod: an RC value of 0.243 was produced, whereas without filters, the RC was only 0.235. With a dual-channel guidance, an even better RC is achieved as it uses anatomical information on tissue boundaries alongside the enhanced input image. When the blurring was slight, it produced an RC of 0.244, which was the highest of the RC values obtained. The Dual-Channel Guided Filter provides



(a) Performing filtering after each iteration (method of sieves)[†]



(b) Performing filtering only after the last iteration (post-reconstruction filtering)

Figure 2: Recovery coefficient (RC) of the thinnest rod of the NEMA phantom as a function of the filters' blur strength, which is characterized by the standard deviation of the uniform region activity. Abbreviations: IGF: Input Guided Filter, EIGF: Enhanced Input Guided Filter, AGF: Anatomically Guided Filter, AGIGF: Anatomically-Guided-Input Guided Filter, DCGF: Dual-Channel Guided Filter.

a better recovery coefficient than any other filter even with increasing blur strength.

4.1.2. Runtime

Table 1 summarizes the runtime of the filters at different filter window sizes. The growing runtime of the median filtering was caused by the sorting needed to determine the median. (Note that constant-time algorithms for median filtering already exist [8, 12].) In contrast, Gaussian filtering can be executed very quickly due to the separability of the operation. Guided filtering has a relatively high runtime, but it is independent of the filter window size. Thus, for large filter windows, it is faster than bilateral filtering.

r	median	Gaussian	bilateral	IGF	EIGF	AGF	AGIGF	DCGF
1	0.0410	0.0089	0.3535	0.4857	0.6030	0.5097	0.6584	1.2466
2	1.0632	0.0090	0.3533	0.4865	0.6048	0.5091	0.6580	1.2488
3	9.2433	0.0092	0.3542	0.4867	0.6034	0.5093	0.6582	1.2464
4	44.4551	0.0094	0.3596	0.4880	0.6041	0.5095	0.6581	1.2462
5	178.0231	0.0096	0.3754	0.4867	0.6034	0.5092	0.6585	1.2468
6	478.2444	0.0098	0.4003	0.4866	0.6052	0.5100	0.6581	1.2465
7	1116.0873	0.0101	0.4322	0.4861	0.6038	0.5095	0.6587	1.2462
8	2308.0518	0.0103	0.4767	0.4859	0.6055	0.5098	0.6582	1.2473
9	4412.3638	0.0105	0.5293	0.4860	0.6044	0.5095	0.6581	1.2517
10	7911.3560	0.0107	0.5968	0.4868	0.6031	0.5097	0.6580	1.2454

Table 1: Filtering time in seconds as a function of the filter window radius, averaged over 100 runs. For guided filter variants, runtime includes the preparation of the guidance. Abbreviations: IGF: Input Guided Filter, EIGF: Enhanced Input Guided Filter, AGF: Anatomically Guided Filter, AGIGF: Anatomically-Guided-Input Guided Filter, DCGF: Dual-Channel Guided Filter.

For the guided filter variants, runtimes shown in Table 1 include the preparation of the guidance as well. The most time is required by the preparation of the anatomically guided input, which involves handling the different resolution of the modalities and performing guided bilateral filtering.

Dual-channel guided filtering has a higher runtime due to matrix operations: it takes an average of 1.0179 seconds to execute filtering, as opposed to the 0.4865 seconds of the single-channel filtering. The preparation time of two guidances leads to a further increase in runtime.

4.1.3. Conclusion

By selecting the appropriate guidance and parameterization, the guided filter is extremely effective in suppressing noise while preserving region boundaries. The best results were obtained with the dual-channel guided filter, which uses information on tissue boundaries provided by an anatomical image, and the enhanced version of the input image as well. However, its higher computational requirements lead to higher runtime, which may not be desirable if the filtering is performed after every iteration. In this case, choosing a single-channel guidance such as the enhanced or the anatomically guided input image should be considered. The best alternative to the guided filter was the bilateral filter, but the latter blurred the fine details of the image slightly more, thus reducing the diagnostic utility of the resulting image. Furthermore, bilateral filtering can occasionally introduce false edges in the image, which does not occur with guided filtering [4].

With the number of iterations examined (30 iterations), there was no significant difference whether filtering was performed after every iteration or only after the last iteration. However, for other data or for longer reconstructions, performing filtering during reconstruction can be necessary to keep the variance of our estimation acceptably low.




Figure 3: Axial slices of the rat phantom

4.2. Noise Reduction in Dynamic Reconstructions

Evaluation was performed on a rat phantom (Fig. 3), which consists of four regions: body, lung, striatum and cerebellum.

4.2.1. Results using the two-tissue compartment model

Fig. 4a shows the image reconstructed without filters using two-tissue compartment model and 50 iterations. A high level of noise can be observed. Even a single application of the guided filter at the end of the reconstruction greatly suppressed this noise (Fig. 4). However, with the wrong filter parameters, artifacts appeared in the output image (Fig. 5).

Fig. 5a shows that if the Triple-Channel Guided Filter is applied after every iteration with low regularization ($\epsilon = 1$), high-variance areas are formed along region boundaries. This is because if the value of ϵ is small, then $\sigma_k^2 \gg \epsilon$ at boundaries, so $a_k \approx 1$ and $b_k \approx 0$ will hold true and therefore the output image will match the guidance. The guidance in this case is multi-channel, where each channel is the sum of the activity images of a frame group. These activity images are noisy in themselves, and noise is cumulative. This accumulated noise is transferred to the output, leading to an even noisier guidance in the next iteration, thus triggering a chain




Figure 4: Reconstruction output without filters, and with activity filtering after the last iteration. Two-tissue compartment model. Abbreviations: AGIGF: Anatomically-Guided-Input Guided Filter, SRGF: Static Reconstruction Guided Filter, TCGF: Triple-Channel Guided Filter.

process that results in worse and worse images. The problem can be eliminated by increasing the regularization to $\epsilon = 10$ (Fig. 6a). Another solution may be given if the anatomical segmentation of the patient is available. In this case, voxels that do not belong to the same tissue as the voxel in the center of the filter window can be ignored during the calculation of the coefficients. If the filtering does not cross region boundaries, much more accurate results are obtained along the edges (Fig. 6b.).

Fig. 5b shows the result of parametric filtering. The guidance was the image of the static reconstruction, but other guidances produced similar results. The cerebellum contains considerable noise and a halo appeared along its outline. The contour of the body was blurred and scattered low-intensity dots formed throughout the image. The reason for this phenomenon is that voxel activity is an exponential function of the parameters of the two-tissue compartment model. The guided filter, similarly to the Gaussian and the bilateral filters, assumes linear scaling and is therefore not applicable to exponentials [15].

4.2.2. Results using the spectral model

Fig. 7 shows the image reconstructed without filters, using the spectral model.

Activity filtering. The Anatomically-Guided-Input Guided Activity Filter produced a very good image even with low regularization (Fig. 8a). In case of low regularization, there




Figure 5: Image artifacts using two-tissue compartment model




Figure 6: Solutions to the image artifact in Fig. 5a

was no visible difference between the outputs when region boundaries were allowed to be crossed and when they were not. However, in case of high regularization, the image became blurred if boundaries were not considered (Fig. 8b).

The Static Reconstruction and the Triple-Channel Guided Filter did not reduce the noise when the regularization was low. However, with high regularization, the smoothing effect was well observable. In this case, the edges in the image were blurred if the region boundaries were allowed to be crossed during filtering (Fig. 10a). However, when boundaries were considered, the edges were preserved extremely accurately (Fig. 10b). The best results were achieved by the Triple-Channel Guided Filter with high regularization and



Figure 7: Reconstruction output without filtering. Spectral model.




Figure 8: Reconstruction output after activity filtering in the last iteration using the anatomically guided input as guidance. Spectral model.




Figure 11: Reconstruction output after activity filtering in every iteration using the triple-channel guidance. Spectral model.




Figure 9: Reconstruction output after activity filtering in the last iteration, with low regularization ($\epsilon = 1$). Spectral model.




Figure 12: Reconstruction output after parametric filtering in the last iteration using the triple-channel guidance and low regularization ($\epsilon = 1$). Spectral model.

with taking into account the region boundaries. The reconstructed time-activity function is shown in Fig. 15.

When filtering was performed after every iteration, all the three examined guided filter variants produced a very smooth image even with low regularization (Fig. 11a). However, with high regularization, the image was blurred if the region boundaries were crossed (Fig. 11b). Best results of the all the reconstructions were produced by applying the Triple-Channel Guided Filter after every iteration with low




Figure 10: Reconstruction output after activity filtering in the last iteration using the triple-channel guidance, with high regularization ($\epsilon = 1000$). Spectral model.

regularization, even without information on region boundaries (Fig. 16).

Parametric filtering. The voxel activity is a linear function of the parameters of the spectral model. Therefore, parametric filtering can be performed without distorting the results.

Parametric filtering resulted in identical images with all three guidances. Nor did the regularization strength affect the output. Prohibiting boundary crossing during calculations proved to be crucial, as without it, the reconstructed image was blurred (Fig. 12a and Fig. 13a). However, when the boundaries were considered, the output (Fig. 12b and Fig. 13b) was just as sharp and noise-free as that of the Anatomically-Guided-Input Guided Activity Filter, and parametric filtering requires fewer filtering operations, since the number of kinetic parameters is usually less than the number of frames. The time-activity function produced is shown in Fig. 17.

4.2.3. Conclusion

The best results were achieved by the Triple-Channel Guided Activity Filter. When filtering was performed only after the last iteration, then high regularization, and when it was performed after every iteration, then low regularization pro-




Figure 13: Reconstruction output after parametric filtering in the last iteration using the triple-channel guidance and low regularization ($\epsilon = 1$). Spectral model.




Figure 14: Time-activity function without filters

duced the better results. In both cases, the output image was almost entirely noise-free.

The sharpness of the image can be further improved if the anatomical segmentation of the patient is available and region boundaries are not crossed during the calculation of the filtering coefficients. (Performing filtering after every iteration will produce a sufficiently sharp image even without this knowledge, but it may be necessary to use this additional information if filtering is performed only once.) Parametric filtering produces a sharp image only if region boundaries are not crossed during calculations, but in that case, it achieves just as good results as activity filtering and requires less executions.

4.3. Partial Volume Correction

Fig. 18a shows the reconstruction output without PVC. It is noticeable that region boundaries were blurred. Performing the proposed PVC technique restores the sharpness of the boundaries (Fig. 18b). The reconstructed time-activity functions are shown in Fig. 19 and Fig. 20. When correction was applied, the average voxel activity was significantly closer to the reference than without correction.




Figure 15: Time-activity function after performing triple-channel-guided activity filtering in the last iteration with high regularization ($\epsilon = 1000$) and knowledge on region boundaries.




Figure 16: Time-activity function after performing triple-channel-guided activity filtering in every iteration with low regularization ($\epsilon = 1$), without considering region boundaries.

5. Summary

In this paper, the guided filter was analyzed for improving PET imaging. First, its noise-reducing capability was examined, in which the filter proved to be particularly effective given the right guidance and parameterization. In addition to suppressing noise, the filter can also be used to sharpen tissue boundaries in the image.

For static reconstructions, a dual-channel guidance was suggested, where the first channel was the enhanced input and the second was the anatomically guided input. For dynamic reconstructions, a triple-channel guidance was recommended, where the channels were obtained by grouping frames and summing the activity images of each group. Using the proposed guidances, guided filtering yielded significantly smoother and sharper images than the alternative filtering methods (median, Gaussian, bilateral). For dynamic PET reconstructions, a new approach called parametric filtering was also conceived, which performs filtering on kinetic parameters. If the anatomical segmentation is known,




Figure 17: Time-activity function after performing anatomically-guided-input-guided parametric filtering in the last iteration with low regularization and knowledge on region boundaries.




Figure 18: Reconstruction output with and without PVC. Spectral model, correction was applied after the last iteration.

parametric filtering produces just as good results as activity filtering, but more efficiently.

In the last section, a guided-filter-based partial volume correction technique was presented, which was able to efficiently compensate for PVE by using the PSF of the system.




Figure 19: Time-activity function without PVC




Figure 20: Time-activity function with PVC ($\epsilon = 1$)

Acknowledgements

The research has been supported by the European Union, co-financed by the European Social Fund (EFOP 3.6.1-16-2016-00014, Research, Development, and Educational Integration of Disruptive Technologies in the Field of E-mobility; EFOP-3.6.2-16-2017-00013, Thematic Fundamental Research Collaborations Grounding Innovation in Informatics and Infocommunications), and by the ÚNKP-19-1 New National Excellence Program of the Ministry for Innovation and Technology.



References

1. CUDA Zone | NVIDIA Developer. <https://developer.nvidia.com/cuda-zone>. Accessed: 2019-10-28.
2. L. Caraffa, J. Tarel, and P. Charbonnier. The guided bilateral filter: When the joint/cross bilateral filter becomes robust. *IEEE Transactions on Image Processing*, 24(4):1199–1208, April 2015.



3. Andrew L. Goertzen et al. Nema nu 4-2008 comparison of preclinical pet imaging systems. *Journal of nuclear medicine : official publication, Society of Nuclear Medicine*, 53(8):1300–1309, Aug 2012.
4. Kaiming He, Jian Sun, and Xiaoou Tang. Guided image filtering. *IEEE transactions on pattern analysis and machine intelligence*, 35:1397–1409, 06 2013.
5. T J Herbert. Statistical stopping criteria for iterative maximum likelihood reconstruction of emission images. *Physics in Medicine and Biology*, 35(9):1221–1232, sep 1990.
6. Hongbo Liu, Kun Wang, and Jie Tian. Postreconstruction filtering of 3d pet images by using weighted higher-order singular value decomposition. *BioMedical Engineering OnLine*, 15, 12 2016.
7. L. Lu, D. Hu, X. Ma, J. Ma, A. Rahmim, and W. Chen. Dynamic pet denoising incorporating a composite image guided filter. In *2014 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–4, Nov 2014.
8. Z. Ma, K. He, Y. Wei, J. Sun, and E. Wu. Constant time weighted median filtering for stereo matching and beyond. In *2013 IEEE International Conference on Computer Vision*, pages 49–56, Dec 2013.
9. M. Magdics, L. Szirmay-Kalos, B. Tóth, and T. Umenhoffer. Filtered sampling for pet. In *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, pages 2509–2514, Oct 2012.
10. J. Nuyts and J. A. Fessler. A penalized-likelihood image reconstruction method for emission tomography, compared to postsmoothed maximum-likelihood with matched spatial resolution. *IEEE Transactions on Medical Imaging*, 22(9):1042–1052, Sep. 2003.
11. László Papp, Gábor Jakab, Balázs Tóth, and László Szirmay-Kalos. Adaptive bilateral filtering for PET. In *IEEE Nuclear science symposium and medical imaging conference, MIC'14*, pages M18–104, 2014.
12. S. Perreault and P. Hebert. Median filtering in constant time. *IEEE Transactions on Image Processing*, 16(9):2389–2394, Sep. 2007.
13. L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122, Oct 1982.
14. Donald L. Snyder and M.I. Miller. The use of sieves to stabilize images produced with the em algorithm for emission tomography. *IEEE Transactions on Nuclear Science*, 32(5):3864–3872, Oct 1985.
15. László Szirmay-Kalos and Ágota Kacsó. Regularizing direct parametric reconstruction for dynamic pet with the method of sieves. In *Molecular Imaging Conference, MIC '16*, pages M16D–1, 2016.
16. László Szirmay-Kalos, Ágota Kacsó, Milán Magdics, and Balázs Tóth. Dynamic pet reconstruction on the gpu. *Periodica Polytechnica Electrical Engineering and Computer Science*, 62(4):134–143, 2018.
17. C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pages 839–846, Jan 1998.
18. Mattia Veronese, Gaia Rizzo, Alessandra Bertoldo, and Federico E. Turkheimer. Spectral analysis of dynamic pet studies: A review of 20 years of method developments and applications. *Computational and Mathematical Methods in Medicine*, 2016(7187541), 2016.
19. Guobao Wang and Jinyi Qi. An optimization transfer algorithm for nonlinear parametric image reconstruction from dynamic pet data. *IEEE Trans Med Imaging*, 31(10):1977–1988, 2012.
20. Hiroshi Watabe, Yoko Ikoma, Yuichi Kimura, Mika Naganawa, and Miho Shidahara. Pet kinetic analysis—compartmental model. *Annals of Nuclear Medicine*, 20(9):583, Nov 2006.
21. Jianhua Yan, Jason Chu-Shern Lim, and David W Townsend. MRI-guided brain PET image filtering and partial volume correction. *Physics in Medicine and Biology*, 60(3):961–976, jan 2015.
22. Daniel Yu and Jeffrey Fessler. Edge-preserving tomographic reconstruction with nonlocal regularization. *Medical Imaging, IEEE Transactions on*, 21:159 – 173, 03 2002.