

Comparison of Non-Linear Filtering Methods for Positron Emission Tomography

Dóra Varnyú and László Szirmay-Kalos

Abstract—As a result of the limited radiotracer dose, acquisition time and scanner sensitivity, positron emission tomography (PET) images suffer from high noise. In the current clinical practice, post-reconstruction filtering has become one of the most common noise reduction techniques. However, the range of existing filters is very wide, and choosing the most suitable filter for a given measurement is far from simple. This paper aims to provide assistance in this choice by comparing the most powerful image denoising filters, covering both image quality and execution time. Emphasis is placed on non-linear techniques due to their ability to preserve edges and fine details more accurately than linear filters. The compared methods include the Gaussian, the bilateral, the guided, the anisotropic diffusion and the non-local means filters, which are examined in both static and dynamic PET reconstructions.

Index Terms—positron emission tomography, image denoising, post-reconstruction filtering, gaussian, bilateral, median, anisotropic diffusion, non-local means, guided filter, efop

I. INTRODUCTION

POSITRON emission tomography (PET) is an imaging technique used to observe biochemical or pharmacological processes in the body. As it provides functional information, PET is a particularly helpful tool for early diagnosis of diseases and pharmacokinetic studies. However, the applicable radioactive dose and the acquisition time are severely limited (cost, physiological effect) and the sensitivity of the imaging system is also generally low. Because of this, PET images suffer from high levels of noise, which can make small lesions such as early-stage tumors impossible to spot.

One possible way to suppress noise is to introduce a penalty term into the maximum-likelihood optimization [1]. However, determining the appropriate parameter values is challenging because they depend on the measured data and the reconstruction settings. Moreover, if the penalty function is not convex, optimization becomes complex and resource-intensive [2].

Another possible solution for noise reduction is early termination [3], which involves ending the iterative reconstruction algorithm well before its convergence. The determination of a stopping point is quite challenging and usually a compromise has to be made between image detailedness and noisiness.

This study focuses on post-reconstruction filtering for image denoising. There are various filters in current clinical application with different characteristics and resource requirements. Due to its simplicity, the Gaussian filter [4] is most commonly used, but it smoothes out image structures such

as tissue boundaries or small lesions besides the noise, thus can deteriorate important clinical information. Better results may be achieved using non-linear filters, since they better preserve the non-linear features of the image such as edges and boundaries. Moreover, certain types of noise can only be effectively removed with non-linear filters. A common example is salt-and-pepper noise, against which the median filter is most effective [4]. Another widely used non-linear filter is the bilateral filter [5], [6], which can also be considered as an extension of the Gaussian kernel. However, in certain scenarios, the bilateral filter can introduce false edges in the image (gradient reversal problem) [7]. An alternative that is free of the gradient reversal problem is the guided filter [7], which produces its output using a guidance image. As the guidance can come from another imaging modality (e.g. CT or MRI) [8], guided filtering makes it possible to take into account anatomical tissue boundaries during filtering. The drawback of this filter is that it is challenging to determine the best guidance and parameter settings as there are many options to choose from and they must be tuned to the measured data. In our previous work [9], we have investigated this topic and proposed several promising guidances for both static and dynamic reconstructions.

Another tool that is able to incorporate high-resolution anatomical images to enhance the output of PET is the anisotropic diffusion (AD) filter [10], [11]. However, it often results in artificially piecewise smooth regions [12]. A more recent alternative is the non-local means (NLM) filter [13], which smooths intensities by the weighted average of intensities in a large neighborhood according to their similarities and was used successfully for PET image denoising in various scenarios [12], [14].

The purpose of this paper is to provide a comprehensive comparison of the different filters for PET in terms of the quality of the output image and the runtime of the operation.

The structure is as follows. In Section II, we give a brief overview of the examined filters. Section III analyzes noise reduction in static reconstructions, while dynamic reconstructions are investigated in Section IV. Finally, the paper is closed with conclusions in Section V.

II. OVERVIEW OF EXAMINED FILTERS

Filtering computes output image Q from input image P by either a linear or a non-linear algorithm. To visually compare the outputs of the examined methods, we performed filtering on a mouse scan measured on Mediso's *nanoScan PET/CT* (Fig. 1) and on a human scan measured on Mediso's *AnyScan human PET/CT* (Fig. 2).




Fig. 1: Post-reconstruction filtering of a mouse scan using different filter algorithms




Fig. 2: Post-reconstruction filtering of a human scan using different filter algorithms

A. Gaussian Filter

The Gaussian filter replaces voxel activities with the Gaussian-weighted average of the activities of adjacent voxels. That is, it convolves with a filter kernel containing Gaussian weights $g_\sigma(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot \exp(-\frac{x^2}{2\sigma^2})$, where x is an arbitrary real number parameter and σ is the standard deviation.

B. Bilateral Filter

The operation of the bilateral filter is similar to that of the Gaussian filter, however, the weights depend not only on the Euclidean distances of the voxels but also on the differences in their activities. The filtering operation is:

$$Q_i = \frac{\sum_{j \in \Omega_i} P_j \cdot g_\sigma(\|i - j\|) \cdot g_\sigma(P_i - P_j)}{\sum_{j \in \Omega_i} g_\sigma(\|i - j\|) \cdot g_\sigma(P_i - P_j)}, \quad (1)$$

where Ω_i is the window centered at voxel i whose activity is currently being calculated, g_σ is the Gaussian function with standard deviation σ , and $\|\cdot\|$ denotes Euclidean distance.

C. Guided Filter

The guided filter [7] is based on a local linear model between a guidance image G and the output image Q :

$$Q_i = a_k G_i + b_k, \quad \forall i \in \Omega_k, \quad (2)$$

where a_k and b_k are coefficients assumed to be constant in the window Ω_k centered at voxel k . Using this model, the output image will be closest to the input if

$$a_k = \frac{\frac{1}{|\Omega|} \sum_{i \in \Omega_k} G_i P_i - \mu_k \bar{P}_k}{\sigma_k^2 + \epsilon}, \quad (3)$$

$$b_k = \bar{P}_k - a_k \mu_k, \quad (4)$$

where ϵ is a regularization constant, 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 .

This way, the output can be calculated as follows:

$$Q_i = \frac{1}{|\Omega|} \sum_{k|i \in \Omega_k} (a_k G_i + b_k) = \bar{a}_i G_i + \bar{b}_i, \quad (5)$$

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 most important task is to choose a guidance so that the outlines of the tissues are kept as sharp as possible while noise is suppressed properly.

One option is to simply use the input image as the guidance. However, because of the local linear relationship between the guidance and the output, an extensively noisy guidance might transfer the noise into the filtered image. If the input is expected to have high noise, denoising might be advantageous before using it as a guidance. In our previous work [9], we applied a Gaussian filter followed by a high-boost filter.

With the spreading of combined PET/CT and PET/MRI scanners, it also becomes possible to use an anatomical image as guidance, thus incorporating tissue boundary information into the filtering. However, since different modalities measure different physical quantities, using their output directly as a guidance can introduce features that are only present in the other modality. To avoid this, we have proposed a joint bilateral filtering algorithm to create a new guidance that uses the anatomical image only indirectly [9]:

$$G_i = \frac{\sum_{j \in \Omega_i} P_i \cdot g(\|i - j\|) \cdot g(A_i - A_j)}{\sum_{j \in \Omega_i} g(\|i - j\|) \cdot g(A_i - A_j)}, \quad (6)$$

where P is the input image, A is the anatomical image, Ω_i is the window centered at voxel i , g is the Gaussian function, and $\|\cdot\|$ denotes the Euclidean distance.

Different guidances can also be combined to form a multi-channel guidance. In the static reconstructions, we worked with a dual-channel guidance whose first channel was the denoised input and the second was the anatomical guidance previously presented. In the dynamic reconstructions, we grouped time frames into three consecutive groups and created a triple-channel guidance by summing the activity images of each of the three frame groups.

D. Median Filter

The median filter replaces the intensity of each voxel with the median of the neighboring voxels. The method is most effective against salt-and-pepper noise, but it can also eliminate low to moderate levels of Gaussian noise.

E. Anisotropic Diffusion Filter

The anisotropic diffusion (AD) filter proposed by Perona and Malik [10] describes an iterative diffusion process

$$P_i^{t+1} = P_i^t + \frac{\partial P_i}{\partial t} = P_i^t + \nabla \cdot (g(|\nabla P|) \nabla P), \quad (7)$$

where t is the time or iteration number, $\nabla \cdot$ is the divergence operator, $|\nabla P|$ is the gradient magnitude, and g is the diffusivity, a non-negative monotonically decreasing function with $g(0) = 1$. In our measurements, we have used diffusivity

$$g(|\nabla P|) = \frac{1}{1 + \frac{1}{2} \left(\frac{|\nabla P|}{K} \right)^3}, \quad (8)$$

where K is a threshold that distinguishes noise from the true signal. We estimate K as a weighted average of gradient magnitudes in a local neighborhood of size $10 \times 10 \times 10$ voxels. Weights are determined by the similarity of the tissue types of the voxels, that is, the difference between their values sampled from an anatomical (CT, MRI) image. This local estimation is then scaled down by a user-defined detail preservation factor (DPF). Increasing the detail preservation factor decreases threshold K above which features are considered as true signal, therefore preserving finer details.

F. Non-Local Means Filter

The non-local means (NLM) filter smooths voxel activities by computing a weighted average of activities in a large search window, with weights determined by the similarity of activities in a smaller local neighborhood (also called patch) of the two voxels being compared:

$$Q_i = \frac{1}{\sum_{j \in \Omega_i} w(i, j)} \sum_{j \in \Omega_i} w(i, j) P_i. \quad (9)$$

In this equation Ω_i is the search window centered at voxel i whose activity is currently being calculated and $w(i, j)$ weight is a measure of similarity between the local neighborhoods (patches) of voxels i and j (denoted by Ψ_i and Ψ_j):

$$w(i, j) = \exp \left(\frac{-\|P(\Psi_i) - P(\Psi_j)\|_{2, \sigma}^2}{h^2} \right), \quad (10)$$

where h is a user-defined smoothing parameter and $\|\cdot\|_{2, \sigma}^2$ stands for the Gaussian-weighted Euclidean distance with $\sigma > 0$ standard deviation of the Gaussian kernel.

III. NOISE REDUCTION IN STATIC RECONSTRUCTIONS

Quantitative comparison of the filters was performed on the NEMA NU 4-2008 preclinical phantom [15] (Fig. 3).




Fig. 3: Slices of the NEMA NU 4-2008 phantom

A. Image quality

We examined the quality of the images produced by the different filtering algorithms based on two metrics: the *recovery coefficient (RC)* and the *contrast-to-noise ratio (CNR)*.

1) *Recovery Coefficients*: In terms of image quality, the most important parts of the NEMA phantom are the five rods of different diameters, particularly the thinnest rod, which is barely visible. Because of its thinness, there is a risk that this rod will disappear as a result of filtering. Therefore, image quality can be characterized by how well the thinnest rod is reconstructed. This is measured by the recovery coefficient (RC), which is the quotient of the reconstructed and the true activity concentration of the rod. When increasing the blur strength of a filter, the RC should not decrease. We describe blur strength by the percentage standard deviation of the activity in the central uniform region (the large contiguous part in the middle, which is well observable in Fig. 3b and Fig. 3c), i.e. the standard deviation divided by the average activity and then multiplied by 100. The RC values of the examined filters as a function of blur strength are plotted in Fig. 4. To achieve different blur strength, filtering parameters (e.g. regularization, Gaussian standard deviation) were changed incrementally.




Fig. 4: Recovery coefficient of the thinnest rod as a function of the percentage standard deviation of the uniform region activity. Abbreviations: IGF: input guided Filter, DCGF: dual-channel guided Filter, NLM51: non-local means filter with search window radius of 5 voxels and patch radius of 1 voxel, NLM75: non-local means filter with search window radius of 7 voxels and patch radius of 5 voxels.

It can be seen that the median filter yielded the worst RC values, but the Gaussian filter also significantly reduced the visibility of the rod. When the blurring was slight, the bilateral filter produced worse RC values than the Gaussian filter, but as the blur strength was increased, the RC decreased to a lesser degree with bilateral filtering.

Compared to these methods, guided filtering resulted in extremely good visibility. Even when simply the input image was used as a guidance, the RC was significantly higher than that of the three previously discussed filters, as it can be clearly seen in in Fig. 4. And with the dual-channel guidance, which incorporates anatomical information on tissue boundaries, an even better RC was achieved. When the blurring was slight, not only did the visibility of the rod not decrease, but it in fact increased: an RC value of 0.2439 was attained, whereas without filters, the RC was only 0.2349.

Anisotropic diffusion filtering resulted in strong blurring at all parameter settings. However, in this range of blur strength, this algorithm gave the highest RC values of all filters.

The non-local means filter was examined in two settings that differed in the size of the search window and the patch window. The smaller version (search window radius of 5 voxels and patch radius of 1 voxel) achieved approximately as good visibility as the dual-channel guided filter, while the bigger version (search window radius of 7 voxels and patch radius of 5 voxels) even outperformed it in the mid-range of blur strength. The highest RC produced by NLM filtering was 0.2354, which is slightly higher than without filtering (0.2349), but still lower than the peak RC value of the dual-channel guided filter (0.2439).

When slight blurring was considered, the dual-channel guided filter, in case of moderate blurring, the non-local means filter, and regarding strong blurring, the anisotropic diffusion filter achieved the best recovery coefficients.

2) *Contrast-To-Noise Ratio*: Another important metric for describing image quality is the contrast-to-noise ratio (CNR), which is calculated as [16]

$$CNR = \frac{\mu_{\text{phantom}} - \mu_{\text{background}}}{\sigma_{\text{background}}}, \quad (11)$$

where μ_{phantom} is the mean activity in the homogeneous phantom, whereas $\mu_{\text{background}}$ and $\sigma_{\text{background}}$ are the mean and the standard deviation of the activity in the background.

Fig. 5 shows the CNR values of the examined filters as a function of blur strength, i.e. the percentage standard deviation of the activity in the central uniform region.




Fig. 5: Contrast-to-noise ratio as a function of the percentage standard deviation of the uniform region activity. Abbreviations: IGF: input guided Filter, DCGF: dual-channel guided Filter, NLM51: non-local means filter with search window radius of 5 voxels and patch radius of 1 voxel, NLM75: non-local means filter with search window radius of 7 voxels and patch radius of 5 voxels.

The Gaussian and the anisotropic filters resulted in very low CNR. However, the bilateral and the median filters, which performed badly regarding the recovery of the thinnest rod, achieved very good contrast-to-noise ratio. The highest

CNR was obtained by the bilateral filter with strong blurring, closely followed by the smaller-window-version of the non-local means filter. The larger-windowed non-local means filter took the lead only when very strong blurring was examined. The dual-channel guided filter also achieved a very good contrast-to-noise ratio – in fact, when the blurring was slight or moderate, it proved to be the best of all filters.

B. Runtime

Runtimes of the Gaussian, the bilateral, the guided, the median, and the non-local mean filters as a function of the filter window radius are summarized in Table I, whereas Table II displays the runtime of the anisotropic filter as a function of the diffusion iteration number.

Gaussian filtering requires negligible time due to the separability of the operation. Bilateral and guided filtering take more time, especially the dual-channel guided filter due to the matrix operations and the preparation time of the anatomical guidance, which involves handling the different resolution of the modalities and performing joint bilateral filtering. However, the execution time of the guided filter is independent of the filter window size. This can make it faster than bilateral filtering for large filter windows.

The increasing runtime of the median filtering was due to the sorting required to determine the median, which was carried out by GPU-based bubble sorting. This is good for small kernels, but for larger ones, other sorting approach should be used instead.

The non-local means filter is very sensitive to both the search window and the path size. If either of the two is large, the runtime can become unacceptably high. Only if the window sizes are small (1-3 voxels) will the runtime be of the same order of magnitude as with bilateral and guided filtering.

Regarding the anisotropic diffusion filter, the execution time is high even if only a few diffusion iterations are carried out. However, it increases linearly with the number of iterations, therefore it can still be faster than the median and the non-local means filters when strong blurring is the goal.

IV. NOISE REDUCTION IN DYNAMIC RECONSTRUCTIONS

In dynamic reconstructions, filtering is applied to the reconstructed (static) images of each time frame in the last iteration.

Quantitative evaluation of the filters was performed on a rat phantom consisting of four homogeneous regions: body, lung, striatum and cerebellum (Fig. 6).




Fig. 6: Slices of the ground truth rat phantom

Fig. 7 shows the reconstruction outputs after performing filtering using the different filter algorithms. Filter parameters were fitted to the measurement data, that is, the best achieved results are presented for each filter. The time-activity functions associated with the images are shown in Fig. 8–Fig. 15. In these graphs, reference activity is indicated by a dashed line, average reconstructed activity by a solid line, and standard deviation of the reconstructed activity by a colored bar around the average (solid) line, where a wider bar means larger standard deviation. Table III presents the quantitative comparison of the filters by displaying the mean and the standard deviation of the reconstructed activity in the striatum (i.e. in the two eye-like region in the output images). The striatum was chosen as the target for the comparison because this is the region where the largest difference can be observed in the filter outputs.

Based on both the output image and the time-activity function, best results were achieved by the non-local means filter with a moderately sized search window and patch window (5 and 1 voxel radius, respectively). It suppressed noise extremely well, as evidenced by the low standard deviation of the reconstructed activity. Boundary edges also remained sharp in all regions. The only place where considerable noise remained is at the edges of the measured volume, which was heavily noisy in the input image due to less data from LORs. This type of noise could only be removed by the median and the anisotropic filters. It should be noted that with a larger search window and patch window, the NLM filter was able to suppress noise better at the volume edges, however, it slightly blurred tissue boundaries.

The second best results were produced by bilateral filtering. Although a few outlier activities can be observed in the output image, the standard deviation of the reconstructed activity is generally low, indicating that most of the noise was suppressed. Furthermore, bilateral filtering preserved tissue boundaries sharp.

Median filtering was not able to completely eliminate the noise, only the outstanding, spike-like values. The less prominent noise values appear as pale dots in the reconstructed image. However, most of the noise at the volume edges was successfully removed. On the other hand, median filtering undesirably reduced the size of the small, but highly active striatum region by replacing voxel activities at the edge of the striatum with the surrounding lower activity of the body.

For guided filtering, we grouped time frames into three consecutive groups and created a triple-channel guidance by summing the activity images of each of the three frame groups. Using this guidance, guided filtering managed to eliminate noise in the body and the cerebellum, but achieved less good results in the lung and the striatum. In addition, it slightly blurred tissue boundaries in some places, such as the right part of the striatum and the nose of the rat.

The anisotropic diffusion filter reduced noise very effectively, obtaining an almost homogeneous activity even at volume edges. However, region boundaries are not as sharp as with other filters, a slight effect can be observed as if tissues had double edges.

Finally, the Gaussian filter was unable to eliminate the noise, only to blur it, compromising region boundaries too.

r	Gaussian	bilateral	IGF	DCGF	median	NLM1	NLM3	NLM5
1	0.0089	0.3535	0.4857	1.2466	0.0410	0.0541	0.6118	2.4043
2	0.0090	0.3533	0.4865	1.2488	1.0632	0.2212	2.9073	11.1695
3	0.0092	0.3542	0.4867	1.2464	9.2433	0.5928	7.9972	30.8299
4	0.0094	0.3596	0.4880	1.2462	44.4551	1.3093	16.9861	65.6742
5	0.0096	0.3754	0.4867	1.2468	178.0231	2.4917	31.0007	120.1560
6	0.0098	0.4003	0.4866	1.2465	478.2444	4.1584	51.1137	192.7670
7	0.0101	0.4322	0.4861	1.2462	1116.0873	6.3966	78.8206	301.5630
8	0.0103	0.4767	0.4859	1.2473	2308.0518	9.3278	115.0860	440.2210
9	0.0105	0.5293	0.4860	1.2517	4412.3638	13.0100	157.4710	616.5830
10	0.0107	0.5968	0.4868	1.2454	7911.3560	17.5618	217.8390	833.3860

TABLE I: Execution time of the filtering algorithms in seconds at different filter window radiuses, averaged over 100 runs. Measurements were run on a NVIDIA TITAN V graphics card [17]. For the dual-channel guided filter, the preparation of the guidance is also included at the indicated time. For non-local means filter, the patch window radius was fixed and the search window radius was changed. Abbreviations: IGF: input guided filter, DCGF: dual-channel guided filter, NLM1, NLM3, NLM5: non-local means filter with patch radius of 1, 3, and 5 voxels, respectively.




Fig. 7: Post-reconstruction filtering of a dynamic rat phantom using different filter algorithms

iterations	anisotropic
20	3.9669
40	5.4686
60	6.9714
80	8.4806
100	9.9869
120	11.4830
140	12.9974
160	14.4686
180	15.9919
200	17.5020

TABLE II: Execution time of the anisotropic diffusion filtering in seconds as a function of the iteration number of the diffusion process, averaged over 100 runs.




Fig. 8: Time-activity function without filters




Fig. 9: Time-activity function after Gaussian filtering with $\sigma = 0.75$ standard deviation




Fig. 10: Time-activity function after bilateral filtering with $\sigma = 2$ spatial standard deviation




Fig. 11: Time-activity function after guided filtering with $\epsilon = 1000$ regularization

	mean	stddev
ground truth	242.7996	0.0000
no filter	155.9013	13.8608
Gaussian	139.8609	19.2666
bilateral	155.2769	9.5501
median	142.8997	27.1341
guided	147.9435	17.8524
anisotropic	133.6422	19.6897
nlm51	155.9950	6.8938
nlm75	155.0453	13.7324

TABLE III: Mean and standard deviation of the activity in the striatum using different filter algorithms, averaged over the entire measurement time




Fig. 12: Time-activity function after median filtering




Fig. 13: Time-activity function after anisotropic filtering with detail preservation factor = 2

V. CONCLUSION

This paper examined post-reconstruction filtering for PET image denoising, comparing the most commonly used filters in terms of image quality and runtime. Quantitative analysis was performed in both a static and a dynamic reconstruction.

In the static reconstruction, the best image quality was achieved by the dual-channel guided filter, followed closely by the non-local means filter. However, both of these filters have a relatively high runtime. When time is important, a single-channel guided filter (e.g. when the guidance is the input image) should be considered.

In the dynamic reconstruction, the best results were obtained by the non-local means filter with a moderately sized search window and patch window. Having slightly more noise, but preserving edges just as sharp, the bilateral filter achieved the second best results. This is also advantageous because bilateral filtering is relatively fast (Table I).

It can be concluded that the best filtering method depends on the measurement data and the reconstruction settings. However, considering both our static and dynamic reconstruction studies, the non-local means filter proved to be the most promising method. In both scenarios, it suppressed noise extremely well and kept tissue boundaries sharp. Its weak point is the high runtime, which should be avoided by setting the search window and the patch window small. In the dynamic reconstruction study, decreasing the window sizes even resulted in a more accurate output image.




Fig. 14: Time-activity function after non-local means filtering with search window radius = 5 voxels, patch radius = 1 voxel, and $h = 5000$ smoothing




Fig. 15: Time-activity function after non-local means filtering with search window radius = 7 voxels, patch radius = 5 voxels, and $h = 2500$ smoothing

ACKNOWLEDGMENT

The research has been supported by the European Union, co-financed by the European Social Fund (EFOP-3.6.2-16-2017-00013, Thematic Fundamental Research Collaborations Grounding Innovation in Informatics and Infocommunications), by the ÚNKP-19-1 New National Excellence Program of the Ministry for Innovation and Technology, and by the OTKA K-124124.

REFERENCES

- [1] 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*, vol. 22, no. 9, pp. 1042–1052, Sep. 2003, doi: 10.1109/TMI.2003.816960.
- [2] D. Yu and J. Fessler, "Edge-Preserving Tomographic Reconstruction with Nonlocal Regularization," *Medical Imaging, IEEE Transactions on*, vol. 21, pp. 159–173, Mar. 2002, doi: 10.1109/42.993134.
- [3] T. J. Herbert, "Statistical stopping criteria for iterative maximum likelihood reconstruction of emission images," *Physics in Medicine and Biology*, vol. 35, no. 9, pp. 1221–1232, Sep. 1990, doi: 10.1088/0031-9155/35/9/003.
- [4] W. K. Pratt, *Digital image processing: PIKS Scientific inside*, 4th ed. Wiley-Interscience, 2007.
- [5] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*. Narosa Publishing House, Jan 1998, pp. 839–846, doi: 10.1109/iccv.1998.710815.

- [6] L. Papp, G. Jakab, B. Tóth, and L. Szirmay-Kalos, "Adaptive bilateral filtering for PET," in *IEEE Nuclear science symposium and medical imaging conference*, ser. MIC'14, 2014, pp. M18–104.
- [7] K. He, J. Sun, and X. Tang, "Guided Image Filtering," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, pp. 1397–1409, Jun. 2013, doi: 10.1109/TPAMI.2012.213.
- [8] J. Yan, J. C.-S. Lim, and D. W. Townsend, "MRI-guided brain PET image filtering and partial volume correction," *Physics in Medicine and Biology*, vol. 60, no. 3, pp. 961–976, Jan. 2015, doi: 10.1088/0031-9155/60/3/961.
- [9] D. Varnyú and L. Szirmay-Kalos, "Guided Filtering and Partial Volume Correction for Positron Emission Tomography," in *Proceedings of the Workshop on the Advances in Information Technology*, ser. Workshop on the Advances in Information Technology (WAIT), B. Kiss and L. Szirmay-Kalos, Eds., no. ISBN 978-963-421-802-9, Budapest, Hungary, Jan. 2020, pp. 89–98.
- [10] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, pp. 629–639, 1990, doi: 10.1109/34.56205.
- [11] L. Szirmay-Kalos, M. Magdics, and B. Tóth, "Volume enhancement with externally controlled anisotropic diffusion," *The Visual Computer*, vol. 33, pp. 331–342, 2017, doi: 10.1007/s00371-015-1203-y.
- [12] W. Qi, T. Xia, X. Niu, C. Ji, M. Winkler, E. Asma, and W. Wang, "A non-local means post-filter with spatially adaptive filtering strength for whole-body PET," in *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. IEEE, Oct. 2015, doi: 10.1109/nssmic.2015.7582060.
- [13] A. Buades, B. Coll, and J.-M. Morel, "A Non-Local Algorithm for Image Denoising," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. IEEE, 2005, doi: 10.1109/cvpr.2005.38.
- [14] C. Chan, R. Fulton, D. D. Feng, and S. Meikle, "Median non-local means filtering for low SNR image denoising: Application to PET with anatomical knowledge," in *IEEE Nuclear Science Symposium Medical Imaging Conference*, Oct. 2010, pp. 3613–3618, doi: 10.1109/NSSMIC.2010.5874485.
- [15] A. L. Goertzen et al., "NEMA NU 4-2008 comparison of preclinical PET imaging systems," *The Journal of Nuclear Medicine*, vol. 53, no. 8, pp. 1300–1309, Aug. 2012, doi: 10.2967/jnumed.111.099382.
- [16] J. Yan, J. Schaefferkoetter, M. Conti, and D. Townsend, "A method to assess image quality for Low-dose PET: analysis of SNR, CNR, bias and image noise," *Cancer Imaging*, vol. 16, p. 26, 2016, doi: 10.1186/s40644-016-0086-0.
- [17] "NVIDIA Titan V GPU specifications," http://www.gpuzoo.com/GPU-NVIDIA/Titan_V.html, accessed on 28/10/2019.



Dóra Varnyú is an MSc student at Budapest University of Technology and Economics. She has contributed to numerous international conferences in the field of medical imaging and computer graphics. Her research focuses on positron emission tomography and has so far covered filtering techniques and motion compensation.



László Szirmay-Kalos was graduated from Budapest University of Technology and Economics in 1987, received Ph.D. and Doctor of Science degree from the Hungarian Academy of Science in 1991 and in 2001, respectively. He is currently a full professor of computer graphics at Budapest University of Technology. His research interests include rendering, Monte Carlo methods and medical imaging. He is the fellow of Eurographics. His web page is <https://www.iit.bme.hu/users/dr-szirmay-kalos-laszlo?language=en>.