

Filtered Sampling for PET

Milán Magdics, László Szirmay-Kalos, Balázs Tóth, Tamás Umenhoffer

Abstract—In tomography reconstruction, the relationship between the finite-element representation of the objective function and the expected number of hits in detectors — or in other words, the particle transport — is described by the system matrix. With the evolution of high-performance hardware, precise on-the-fly estimation of the system matrix becomes more and more feasible, which allows the use of patient-dependent data and makes it unnecessary to deal with the compression of enormous matrices. On-the-fly system matrix generation requires the online approximation of high dimensional integrals, which is usually attacked by Monte Carlo quadrature and importance sampling. Determining the number of samples used by the estimators belongs to the classical tradeoff problem between accuracy and computational time. However, the approximation error mainly comes from the measurement noise and high frequency components of the measured object that cannot be captured using a given sample density. In this paper, we propose the application of filtered sampling for the forward projection step of iterative ML-EM based PET reconstruction to decrease the variance of the integrand and thus to reduce the error of integral estimation for a given set of samples. The input of the forward projection is filtered using a low-pass filter, which reduces noise and increases the probability that samples do not miss high frequency peaks — e.g. a point source. The iteration thus converges to a modified fixed point, from which the original function can be extracted by applying the same filter. The presented model is built into the TeraTomo™ system.

I. INTRODUCTION

In iterative positron emission tomography (PET) forward and back projections alternate. Forward projection models the physics of the system by computing the expected number of simultaneous γ -photon hits in detector pairs (a.k.a. Line Of Responses or LORs), \tilde{y}_L , from the current estimation of the radiotracer density $x(\vec{v})$, while back projection corrects the current estimation based on the ratio of the measured y_L and computed LOR-hits \tilde{y}_L .

The expected hits for a given LOR L is:

$$\tilde{y}_L = \int_{\mathcal{V}} x(\vec{v})T(\vec{v} \rightarrow L)d\vec{v} \quad (1)$$

where \mathcal{V} is the volume of interest, $x(\vec{v})$ is the radiotracer density function and $T(\vec{v} \rightarrow L)$ is the *system sensitivity* denoting the probability that a positron born in \vec{v} causes a γ -photon pair hit in LOR L .

The output of the reconstruction method is the *tracer density* function $x(\vec{v})$, which is approximated in a *finite function series*

This work has been supported by the TeraTomo project of the National Office for Research and Technology, TÁMOP-4.2.2.B-10/1-2010-0009, and OTKA K-719922.

M. Magdics, L. Szirmay-Kalos, B. Tóth, and T. Umenhoffer are with Budapest University of Technology and Economics (e-mail: szirmay@iit.bme.hu).

form:

$$x(\vec{v}) = \sum_{V=1}^{N_{\text{voxel}}} x_V b_V(\vec{v}), \quad (2)$$

where $\mathbf{x} = (x_1, x_2, \dots, x_{N_{\text{voxel}}})$ are unknown coefficients and $b_V(\vec{v})$ ($V = 1, \dots, N_{\text{voxel}}$) are *basis functions*, which are typically defined on a *voxel grid*.

The application of finite elements replaces the continuous problem by a discrete one:

$$\tilde{y}_L = \int_{\mathcal{V}} x(\vec{v})T(\vec{v} \rightarrow L)d\vec{v} = \sum_{V=1}^{N_{\text{voxel}}} \mathbf{A}_{LV} x_V \quad (3)$$

where \mathbf{A}_{LV} is the *system matrix*:

$$\mathbf{A}_{LV} = \int_{\mathcal{V}} b_V(\vec{v})T(\vec{v} \rightarrow L)d\vec{v}. \quad (4)$$

The ML-EM scheme searches tracer density coefficients $x_1, \dots, x_{N_{\text{voxel}}}$ that maximize the probability of measurement results $y_1, \dots, y_{N_{\text{LOR}}}$ by an iterative algorithm started from a uniform tracer density or an initial estimate. Taking into account that the measured hits follow a Poisson distribution, after each forward projection, the ML-EM scheme executes a *back projection* correcting the voxel estimates $x_V^{(n)}$ of iteration step n by scaling factors s_V obtained from the ratios of measured and computed LOR values [9]:

$$x_V^{(n+1)} = x_V^{(n)} \cdot s_V, \quad \text{where } s_V = \frac{\sum_L \mathbf{A}_{LV} \frac{y_L}{\tilde{y}_L}}{\sum_L \mathbf{A}_{LV}}. \quad (5)$$




Fig. 1. The reconstruction as a control loop. Forward projection \mathcal{F} takes the actual voxel values $x_V^{(n)}$ and computes the expectation of LOR events \tilde{y}_L . Back projection \mathcal{B} calculates a correction ratio s_V for every voxel from the expected LOR events \tilde{y}_L and the measured LOR hits y_L .

The reconstruction process can also be interpreted as a control loop (Fig. 1), including forward projection

$$\tilde{y}_L = \mathcal{F}(x) = \sum_{V=1}^{N_{\text{voxel}}} \mathbf{A}_{LV} x_V$$

and back projection

$$s_V = \mathcal{B}(\tilde{y}_L) = \frac{\sum_L \mathbf{A}_{LV} \frac{y_L}{\tilde{y}_L}}{\sum_L \mathbf{A}_{LV}}.$$

This loop is stabilized when $x^{(n+1)} = x^{(n)}$, that is when scaling factors s_V are 1, which means that this loop solves the following equation for x :

$$\mathcal{B}(\mathcal{F}(x)) = 1. \quad (6)$$

System sensitivity $T(\vec{v} \rightarrow L)$ and consequently system matrix elements \mathbf{A}_{LV} are high dimensional integrals since they aggregate all particle paths started at the positron generation and connecting the sequence of particle–matter interaction points by line segments. In the simplest physical model, where positron range and photon scattering are ignored, and detector crystals are assumed to be ideally black, the expected number of hits in LOR L is:

$$\tilde{y}_L = \int_{\mathcal{V}} \int_{\Omega} x(\vec{v}) T(\vec{v}, \vec{\omega} \rightarrow L) d\vec{v} \frac{d\omega}{2\pi}$$

where Ω is the directional hemisphere, and $T(\vec{v}, \vec{\omega} \rightarrow L)$ is 1 if the line of place vector \vec{v} and direction ω intersects the two detector surfaces and zero otherwise. The position of annihilation \vec{v} can be defined by three scalar coordinates, while the direction by two, making this integral five dimensional.

If positron range and photon scattering are also considered, every potential particle–matter interaction point increases the dimension of the integration domain by three, quickly leading to very high dimensional integrals. A classic method to approximate high dimensional integrals is the Monte Carlo quadrature [11]. It has the advantage that the number of its sample points required to reduce the integration error below a threshold does not grow with the dimension of the integration domain unlike in classical deterministic quadrature rules where the required sample number grows exponentially with the dimension of the domain. The computation time of Monte Carlo quadrature is proportional to the number of samples, so it should be as small as possible.

For a given number of samples, the error of Monte Carlo quadrature depends on the distribution of sample points, and the variation of the integrand divided by the sample density. To obtain a better sample distribution, different methods such as importance sampling and low discrepancy sampling were proposed [12]. In the followings, we propose a method to reduce the variation of the integrand.

II. METHODS

Filtered sampling [1] replaces the integrand by another function that has a similar integral but smaller variation, then its integral can be estimated more precisely from discrete samples (Fig. 2). Reducing the variation means the filtering of high frequency fluctuations by a low-pass filter. This filter should eliminate frequencies beyond the limit corresponding to the density of the sample points. On the other hand, it should only minimally modify the integral.

To see how this pre-filtering affects the reconstruction, let us consider the control loop of ML-EM iteration (Fig. 1) and include filtering operator \mathcal{G} before forward projection \mathcal{F} (Fig. 3). Filtering operator \mathcal{G} maps the iteration result x_V to filtered voxel value \hat{x}_V . The modified system stabilizes when the scaling factors s_V are 1, thus we get

$$s_V = \mathcal{B}(\mathcal{F}(\hat{x})) = \mathcal{B}(\mathcal{F}(\mathcal{G}(x))) = 1.$$




Fig. 2. Filtered sampling reduces the approximation error of the integral quadrature by reducing the variance of the integrand.




Fig. 3. The reconstruction loop of filtered sampling. Forward-projection \mathcal{F} computes the expected LOR hits \tilde{y}_L from the filtered voxel values \hat{x}_V that are computed as by applying filter \mathcal{G} to the result of previous iteration $x^{(n)}$. Back projector \mathcal{B} calculates the scaling factor s_V for each voxel, i.e. it obtains the product of the ratios y_L/\tilde{y}_L of the measured and computed LOR hits with the transpose of the system matrix and divides the results by the column sums of the system matrix.

Note that this is the same equation for \hat{x} as the original one (equation 6) for x , thus considering \hat{x} to be the output of the control system, the modified system behaves similarly to the original one. In the modified system we always have two tracer density estimates x_V and \hat{x}_V , that are related as $\hat{x} = \mathcal{G}(x)$. In addition to solution \hat{x} , we also get a sharpened reconstruction $x = \mathcal{G}^{-1}(\hat{x})$.

Note that filtering only affects the input of the forward projection step, the correction made by back projection is applied to the unfiltered estimate of the radiotracer density $x^{(n)}$ (see Fig. 3).

As a low pass filter \mathcal{G} , we experimented with the *Gaussian* and the *Bilateral filters* [14]. The advantage of the Gaussian filter is that it can be defined by its mean and standard deviation, and the mean is conveniently set to zero while the standard deviation is set according to the noise that needs to be suppressed. However, as shown in Fig. 10 and Fig. 11,




Fig. 4. Effect of different filters on a noisy edge. In addition to damping high frequency noise, Gaussian filter also introduces blurring. Bilateral filter, on the other hand, filters noise while keeping sharp transitions of the data.

Gaussian filtering thickens object boundaries. A possible explanation is that edges are not bandlimited signals and have high frequency components beyond the Nyquist limit and the reasonable range of the numerical precision. These high frequency components are eliminated by the Gaussian but cannot be reproduced by its inverse.

Bilateral filters, on the other hand, preserve edges and object boundaries when their parameters are appropriately set. In the most commonly used case, the weights of these filters are products of two Gaussians: one is defined in the spatial domain, the other is in the intensity domain. More specifically, Bilateral filter B is defined as

$$B(\hat{v}, \sigma_d, \sigma_r) = \frac{\int G_{\sigma_d}(\|\vec{v} - \hat{v}\|)G_{\sigma_r}(x(\vec{v}) - x(\hat{v}))x(\vec{v})d\vec{v}}{\int G_{\sigma_d}(\|\vec{v} - \hat{v}\|)G_{\sigma_r}(x(\vec{v}) - x(\hat{v}))d\vec{v}}$$

with \hat{v} denoting the center voxel of the neighborhood and G_σ denoting the one dimensional Gaussian function of standard deviation σ :

$$G_\sigma(t) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{t^2}{2\sigma^2}\right).$$

The intensity-dependent Gaussian weight ensures that neighboring voxels placed on the same side of a step-like signal as the centering voxel \hat{v} get higher weights while voxels from the other side of the edge give less contribution to the filter output, better preserving the edge. The amount of blur is controlled by spatial variance parameter σ_d , while the amount of detail kept is determined by intensity variance parameter σ_r . However, the appropriate value of σ_r is less straightforward to find, since it is given in intensity space which is object dependent.

The optimal spatial parameter for both filters can be obtained from the probability density function (PDF) of the samples used by forward projection (Fig. 5), thus, the filter size may vary through the reconstructed volume. Forward projector operators can be classified into *voxel driven* (a.k.a input driven or direct Monte Carlo) or *LOR driven* (a.k.a output driven or adjoint Monte Carlo) algorithms. Voxel driven methods in most cases sample the volume proportionally to the intensity, which directly generates the PDF in each iteration. LOR driven approaches, on the other hand, usually aim for uniform sample density in LOR-space (a.k.a sinogram-space or image-space), which means that the sample density in voxel-space is approximately the same in every iteration, and the corresponding PDF can be approximated in a preprocessing step.

MIP-mapping can substantially improve the performance of filters having voxel-dependent kernel size on a parallel

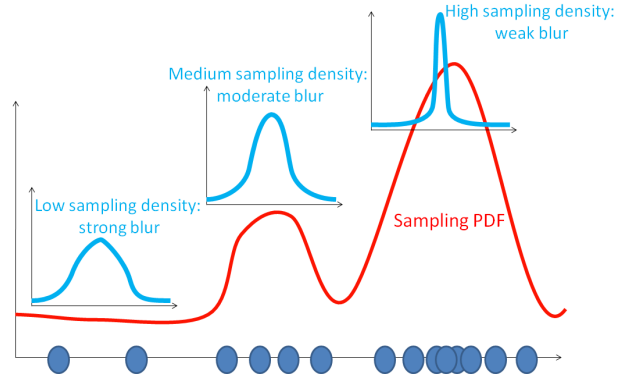


Fig. 5. Spatially varying filtering based on the sampling PDF. Low sampling density cannot capture high variance details, thus, a strong blur is used to decrease the variance of the integrand. When the sampling density is high enough, there is no need to eliminate high frequency details.




Fig. 6. 3D Gaussian-pyramid and MIP-mapping. Lower resolution variants of the volume are created by gradually down-sampling with a Gaussian filter. Spatially varying filtering is executed by sampling the resolution level corresponding to the required amount of blur.

architecture such as the GPU [1]. The first step of MIP-mapping is to build a Gaussian pyramid [8], that is, gradual down-sampling of the original three dimensional voxel array by a factor of 2, by iteratively applying the Gaussian filter (Fig. 6). The upper levels of the pyramid thus correspond to the lower resolution variant of the original volume where the higher frequency details are properly eliminated. After this pre-processing, MIP-mapping executes spatially varying filtering by sampling the volume of the resolution level corresponding to the required level of blurring (Fig. 6). The MIP-map level can be a non-integer scalar pointing between two neighbouring levels of the pyramid. In this case, the final output is the interpolated value of the samples taken from these two levels.

The optimal filter size may be smaller than the voxel size meaning that the local sampling density can already capture

the highest frequency details in the data allowed by the finite element representation. As a consequence, if the time budget given for the reconstruction allows a sufficient sample density (e.g. the maximal distance between samples is smaller than the voxel size throughout the volume of interest), filtered sampling provides no benefits.

III. RECONSTRUCTION METHOD

Our reconstruction method is an OSEM/ML-EM based adjoint Monte Carlo approach. We favor LOR driven, i.e. output driven forward projections since output driven algorithms are more suitable for GPU implementation. In the case of small animal PET where scattering in the material is of less importance, the sampling density of our forward projector is mainly determined by the geometric phase estimating the direct contribution. Thus, in the followings, we describe this component in details. It will be clarified that the algorithm has nearly uniform sampling density in LOR space, making a point-source the worst-case input by wasting most of the samples on nearly zero intensities.

Considering only the geometry, a LOR can be affected only if its detectors are seen at directions $\vec{\omega}$ and $-\vec{\omega}$ from emission point \vec{v} . It also means that emission point \vec{v} and direction $\vec{\omega}$ unambiguously identify detector hit points \vec{z}_1 and \vec{z}_2 , or alternatively, from detector hit points \vec{z}_1 and \vec{z}_2 , we can determine those emission points \vec{v} and direction $\vec{\omega}$, which can contribute.

To establish a LOR-driven approach, we modify our view point from the emission points and directions to detector points, and using the correspondence between them, the detector response is expressed as an integral over the detector surfaces.

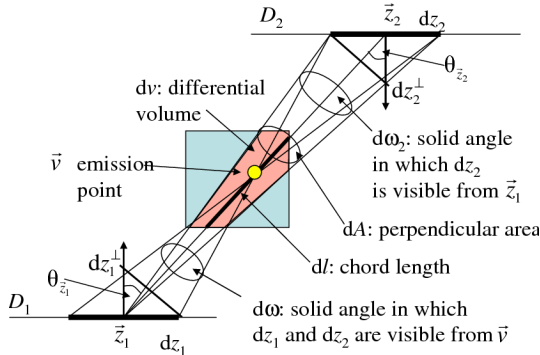


Fig. 7. Computation of the Jacobian of the change of variables. The differential solid angle at which dz_1 detector surface and dz_2 detector surface are simultaneously seen from emission point \vec{v} is $d\omega = d\omega_1 \cos \theta_{z_1} / |\vec{z}_1 - \vec{v}|^2 = d\omega_2 \cos \theta_{z_2} / |\vec{z}_2 - \vec{v}|^2$. The differential solid angle at which dz_2 is seen from point \vec{z}_1 is $d\omega_2 = d\omega_2 \cos \theta_{z_2} / |\vec{z}_2 - \vec{z}_1|^2 = dA / |\vec{z}_1 - \vec{v}|^2$. Finally, the differential volume intersected by lines of \vec{z}_1 and \vec{z}_2 is $dv = dl dA$, where dl is the length of the line segment intersecting dv , and dA is the surface area that is perpendicular to the line.

According to Fig. 7, the Jacobian of the change of integration variables is:

$$d\omega dv = \frac{\cos \theta_{z_1} \cos \theta_{z_2}}{|\vec{z}_1 - \vec{z}_2|^2} dl dz_1 dz_2$$

where θ_{z_1} and θ_{z_2} are the angles between the surface normals and the line connecting points \vec{z}_1 and \vec{z}_2 on the two detectors,

respectively. With this, the LOR integral can be expressed as a triple integral over the two detector surfaces D_1 and D_2 of the given LOR and over the line connecting two points \vec{z}_1 and \vec{z}_2 belonging to the two detectors:

$$\tilde{y}_L = \int_{D_1} \int_{D_2} \frac{\cos \theta_{z_1} \cos \theta_{z_2}}{2\pi |\vec{z}_1 - \vec{z}_2|^2} \left(\int_{\vec{z}_1}^{\vec{z}_2} x(\vec{l}) dl \right) dz_2 dz_1. \quad (7)$$

LOR driven methods are also called *ray based* since they identify voxels that may contribute to a LOR by casting one or more rays between two points on the LOR detectors. For the approximation of the line integral along a ray, we may use ray marching, Siddon's algorithm [10], Joseph's method [3], or distance based techniques [7]. Siddon's method and Joseph's method assume piece-wise constant basis functions, which can be extended to piece-wise tri-linear basis [13]. Ray marching, on the other hand, is appropriate for any basis function selection. Ray marching takes constant length steps along the ray and samples the activity and the material density at finite number of sample points. The first sample location is jittered randomly to guarantee unbiased estimate for the line integral.

Equation (7) can be estimated by taking N_{detline} uniformly distributed point pairs, $(\vec{z}_1^{(i)}, \vec{z}_2^{(i)})$ on the two detectors, and selecting N_{march} equidistant points \vec{l}_{ij} along each line segment $(\vec{z}_1^{(i)}, \vec{z}_2^{(i)})$ (Fig. 7):

$$\tilde{y}_L \approx \frac{D_1 D_2}{2\pi N_{\text{detline}}} \sum_{i=1}^{N_{\text{detline}}} \left(\frac{\cos \theta_{z_1^{(i)}} \cos \theta_{z_2^{(i)}}}{|\vec{z}_1^{(i)} - \vec{z}_2^{(i)}|^2} \sum_{j=1}^{N_{\text{march}}} x(\vec{l}_{ij}) \Delta l_i \right).$$

This formula is the Monte Carlo estimator of the expected LOR hits taking discrete point samples \vec{l}_{ij} in the voxel domain. In order to get a high accuracy estimate, the domain of each basis function that is relevant for this LOR should be sufficiently densely sampled, which would lead to very high sample numbers. However, if the average number of samples in each basis function is about one or less, this estimate has non-negligible variance especially if the activity has high frequency variations. The proposed filtering scheme can be used to smooth these variations before the Monte Carlo sampling.

IV. RESULTS

The proposed method is integrated into the TeraTomoTM system [4]. The filtering step has been implemented in CUDA [6] and it has been executed on an NVIDIA GeForce 480 GFX GPU. Due to the high arithmetic performance and bandwidth of the GPU, the execution time of the filtering step is negligible comparing to the projection operators even for higher resolution volumes. Thus, our proposed method has practically no overhead.

To demonstrate the positive effects of using filtered sampling, we modeled Mediso's *nanoScan-PET/CT* [5], which has 12 detector modules consisting of 81×39 crystals of size $1.12^2 \times 13 \text{ mm}^3$ and simulated a noisy measurement of a Derenzo-like phantom where the simulation corresponds to a 10 second measurement, and a point-source phantom with

GATE [2]. In both cases, we reconstructed the measured values with and without filtered sampling. We used only $N_{\text{detline}} = 1$ and $N_{\text{march}} = 36$ samples in both cases and set standard deviation σ of the Gaussian to 2 (Fig. 10). Note that by using the presented filtering method, approximately the same image quality could be achieved as obtained with 8 times more samples by the original method.




Fig. 8. Cross Correlation errors for different filtering kernel widths when the Derenzo is reconstructed at double, i.e. $288 \times 288 \times 256$ resolution




Fig. 9. The effect of the filtering between forward and back projections using different filter kernels when the Derenzo is reconstructed at double, i.e. $288 \times 288 \times 256$ resolution. We computed 100 iterations. As the no-filtering and filtering with $\sigma = 0.3$ reconstructions diverge, their images are shown after 60 iterations. The $\sigma = 0.5$ reconstruction also deteriorates and should be scaled by a factor of 6.

To demonstrate that filtering allows to increase the resolution without the need of significantly more samples, we reconstructed a higher activity Derenzo at $288 \times 288 \times 256$ resolution as well without increasing the number of samples




Fig. 11. Line profile for the Derenzo phantom of Fig. 10. Note that filtered sampling using the Gaussian filter makes the rods narrower while the bilateral filter better preserves their thickness.




Fig. 12. Reconstruction of the linearity-homogeneity phantom simulated by GATE.

(Fig. 8 and Fig. 9). At such high resolution, the Gaussian clearly demonstrates its potential since without using it, the iteration process does not converge.

Fig. 12 and Fig. 13 show the results for the GATE-simulated linearity homogeneity phantom (corresponding to a 2 second measurement), consisting of eight homogeneous cubes having different activity levels. Gaussian filtering greatly reduces the noise coming from the low-dose simulation and the low sampling rate while also stabilizing convergence.

V. CONCLUSION

This paper proposed a filtering method to decrease the variance of the integrand of the high dimensional integrals in the forward projection step of an iterative ML-EM algorithm. We proposed the application of low-pass filtering before the forward projections, while back projection still corrects the original unfiltered voxel array. We have proven that this approach does not compromise the reconstruction and preserves the stability even if high resolution voxel arrays are reconstructed with a low number of Monte Carlo sampling. All steps are implemented on the GPU where the added computational cost of filtering is negligible with respect to forward and back projection calculations.




Fig. 10. Effects of filtered sampling. The upper row of images shows the reconstructions of a Derenzo-like phantom and the bottom row shows the reconstructions of a point source. Note that using filtered sampling results in a similar image quality to a reconstruction using 8 times more samples. The hole artifact of the point source reconstruction in II. a.) arises from the low sampling density, forward projection misses the neighborhood of the point source in several iterations which has a lower probability if the point source is blurred by filtered sampling. Also note that by using Gaussian filter, the boundaries of the object are somewhat thickened, which is not present for the bilateral filter.




Fig. 13. Error curves of the reconstruction of the linearity-homogeneity phantom simulated by GATE.

REFERENCES

- [1] Mark Colbert and Jaroslav Křivánek. Real-time shading with filtered importance sampling. In *ACM SIGGRAPH 2007 sketches*, SIGGRAPH '07, New York, NY, USA, 2007. ACM.
- [2] GATE. GATE, the Geant4 application for emission tomography, 2011. Available: <http://opengatecollaboration.healthgrid.org/> [January 1 2011].
- [3] Peter M. Joseph. An improved algorithm for reprojecting rays through pixel images. *IEEE Transactions on Medical Imaging*, 1(3):192–196, nov. 1982.
- [4] M. Magdics, L. Szirmay-Kalos, B. Toth, D. Legrady, A. Cserkaszy, L. Balkay, B. Domonkos, D. Volgyes, G. Patay, P. Major, J. Lantos, and T. Bukki. Performance Evaluation of Scatter Modeling of the GPU-based “Tera-Tomo”; 3d pet reconstruction. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE*, pages 4086–4088, oct. 2011.
- [5] Mediso Medical Imaging Systems. title = NanoScan-PET/CT, <http://www.mediso.com/products.php?fid=2,11&pid=86>
- [6] Nvidia. CUDA, Compute Unified Device Architecture, 2011. Available: http://www.nvidia.com/object/cuda_home.html [January 1 2011].
- [7] A. J. Reader and H. Zaidi. Advances in PET image reconstruction. *PET Clinics*, 2(2):173–190, 2007.
- [8] A. Rosenfeld. *Multiresolution Image Processing and Analysis*. Springer series in information sciences. Springer-Verlag, 1984.
- [9] L. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imaging*, 1:113–122, 1982.
- [10] R. L. Siddon. Fast calculation of the exact radiological path for a three-dimensional ct array. *Medical Physics*, 12(2):252–257, 1985.
- [11] L. Szirmay-Kalos. *Monte-Carlo Methods in Global Illumination — Photo-realistic Rendering with Randomization*. VDM, Verlag Dr. Müller, Saarbrücken, 2008.
- [12] L. Szirmay-Kalos, B. Tóth, M. Magdics, D. Légrády, and A. Penzov. Gamma photon transport on the GPU for PET. *Lecture Notes in Computer Science*, 5910:433–440, 2010.
- [13] László Szirmay-Kalos, Balázs Tóth, and Milán Magdics. Free Path Sampling in High Resolution Inhomogeneous Participating Media. *Computer Graphics Forum*, 30(1):85–97, 2011.
- [14] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proceedings of the Sixth International Conference on Computer Vision, ICCV '98*, pages 839–, Washington, DC, USA, 1998. IEEE Computer Society.