

Variance Stabilizing Transformations in Patch-based Bilateral Filters for Poisson Noise Image Denoising

Arnaud de Decker, John Aldo Lee and Michel Verleysen

Abstract—Denoising is a key step in the processing of medical images. It aims at improving both the interpretability and visual aspect of the images. Yet, designing a robust and efficient denoising tool remains an unsolved challenge and a specific issue concerns the noise model. Many filters typically assume that noise is additive and Gaussian, with uniform variance. In contrast, noise in medical images often has more complex properties. This paper considers images with Poissonian noise and the patch-based bilateral filters, that is, filters that involve a tonal kernel and pair wise comparisons between shifted blocks of the images. The main aim is then to integrate two variance stabilizing transformations that allow the filters to work with Gaussianized noise. The performances of these filters are compared to those of the classical bilateral filter with the same transformations. The experiments include an artificial benchmark as well as a positron emission tomography image.

I. INTRODUCTION

Dealing with noisy data turns out to be one of the toughest challenges in image processing. Noise naturally arises in difficult conditions such as poorly illuminated environments, short exposure times, and low-efficiency photon detectors. The statistical properties of noise will depend on these conditions as well as on the characteristics of the imaging device. In this context, denoising can help recover the underlying signal. In biomedical applications, quality requirements are high and optimal denoising increases both the readability and interpretability of the images. A typical example occurs in radiation oncology, where positron emission tomography (PET) is used to diagnose tumors. Noise and resolution limit the accuracy of the tumor delineation, which can reduce the treatment benefit. This motivates the use of denoising tools that satisfy two conditions. First, they must deal with multiplicative noise, such as Poissonian noise, which is typical in PET data. Second, they have to be edge-preserving, that is, they must be able to attenuate noise without decreasing the resolution and smoothing the underlying signal.

The issue of unsupervised edge-preserving denoising can be addressed with various tools, such as wavelets [1], [2], partial differential equations [3], total variation [4], Bayesian denoising [5], kernel regression [6], gradient approximation [7]. Yet another approach focuses on local filters with a tonal kernel [8]. The local M-smoother (LMS) [9] and the

A. de Decker is with the Machine Learning Group of the Université catholique de Louvain (UCL), Belgium; he is funded by a Belgian F.R.I.A. grant. arnaud.dedecker@uclouvain.be

J.A. Lee is with the Center for Molecular Imaging and Experimental Radiotherapy of the UCL; he is a Postdoctoral Researcher funded by the Belgian National Fund of Scientific Research (FNRS). john.lee@uclouvain.be

M. Verleysen is with the Machine Learning Group of the UCL. michel.verleysen@uclouvain.be

bilateral filter (BF) [10], [11] fall in this category. These filters are generalized in [12] in such a way that the tonal kernel compares image blocks (patches) instead of single pixels. Patches in the tonal kernel were also used outside of the bilateral filter framework by [13], [14].

In the case of Gaussian noise, the patch-based bilateral filter (PBBF) has been shown to outperform both the BF and LMS. However, all filters with a tonal kernel implicitly assume that noise is Gaussian and additive. In order to optimally perform in the case of Poissonian noise, they can be combined with variance stabilizing transformations (VSTs) such as Anscombe's and Fisz's [15]. This paper shows how to integrate the VSTs into both the BF and PBBF. Experiments with an artificial benchmark and a real PET image compare the filters and illustrate the performance gains brought by the VSTs.

The rest of this paper is organized as follows. Section II defines the VSTs, introduces the PBBF, and explains how the VSTs can be integrated into it. Section III describes the experiments and discusses their results. Finally, Section IV gathers the conclusions and some perspectives for future work.

II. THEORY

A. Variance Stabilizing Transformation

A transformation that stabilizes the variance is basically intended to 'Gaussianize' non-Gaussian data. For a given statistical law, an optimal VST converts any sufficiently large sample into another sample that is as close as possible to normality. In the case of Poisson distributions, two well known VSTs are the Anscombe's and Fisz's [15]. If $X \sim \text{Pois}(\lambda)$, Anscombe's VST is defined as $Y = 2\sqrt{X + 3/8}$, where Y is approximately Gaussian with zero mean and unit variance, provided $\lambda > 5$ [16].

Fisz's VST is derived from Fisz's theorem [15], which states that if two independent vectors $X_i \sim \text{Pois}(\lambda_i)$, $i = 1, 2$, then function $\zeta : \mathfrak{R}^2 \rightarrow \mathfrak{R}$ defined as

$$\zeta(X_1, X_2) \doteq \begin{cases} 0 & \text{if } X_1 = X_2 = 0 \\ \frac{X_1 - X_2}{\sqrt{X_1 + X_2}} & \text{otherwise} \end{cases} \quad (1)$$

is such that if $(\lambda_1, \lambda_2) \rightarrow (\infty, \infty)$ and $\lambda_1/\lambda_2 \rightarrow 1$, then

$$\zeta(X_1, X_2) - \zeta(\lambda_1, \lambda_2) \rightarrow N(0, 1) \quad (2)$$

Function ζ is called Fisz's VST. As demonstrated in [16], Fisz's VST converges to normality much faster than Anscombe's VST and is already effective for small values of λ_1 and λ_2 .

B. The Bilateral Filter

Let us consider a D -dimensional image wherein the location of the i th pixel is given by coordinate vector $\mathbf{x}_i = [x_{i1}, \dots, x_{iD}]^T$. The intensity of the i th pixel is written as $f_i = u_i + \varepsilon_i$, where u_i is the noise free signal and ε_i the noise component. In the following developments, we assume that each pixel corresponds to a photon count, and therefore $\varepsilon_i \sim \text{Pois}(u_i)$.

Denosing basically aims at recovering signal u_i starting from noisy measurement f_i . In order to achieve this, the BF replaces intensity f_i with a weighted average of neighboring pixel intensities f_j . The iterative application of this principle yields an estimate of u_i that can be written as

$$\hat{u}_i^{k+1} = \frac{\sum_{j \in N_i} w_{ij} \Psi' \left((\hat{u}_i^k - \hat{u}_j^k)^2 \right) \hat{u}_j^k}{\sum_{j \in N_i} w_{ij} \Psi' \left((\hat{u}_i^k - \hat{u}_j^k)^2 \right)}, \quad (3)$$

In this update, k is the iteration index and N_i denotes the neighborhood of the i th pixel. Symbol w_{ij} is defined as

$$w_{ij} = \exp \left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\rho^2} \right) \quad (4)$$

and modulates the weight of the j th neighbor in the computation of \hat{u}_i^{k+1} ; w_{ij} decays as the distance between \mathbf{x}_i and \mathbf{x}_j grows. In contrast, tonal kernel Ψ' depends on the pixel intensities and is the first derivative of $\Psi(v) = \exp(-\frac{v}{2\sigma^2})$. The BF is a state-of-the-art unsupervised edge-preserving filter, with only two parameters, namely ρ and σ , which are the respective widths of the spatial and tonal kernels. More details about the derivation of the BF can be found in [8] and [17].

C. Patch-Based Bilateral Filter

The PBBF [12] extends the classical bilateral filter in order to use patches of the image instead of single pixels in the tonal kernel. This generalization has been shown to significantly improve the filtering performances [12]. Considering an iterative filter, patch P_i^k refers to the vector of intensities that are found at the k th iteration in the image block centered on \mathbf{x}_i . The block size is equal to p^D , where p is a parameter of the PBBF, constrained to be an odd integer. Patches can be compared, for instance by using the Euclidean norm. The distance between P_j^k and P_i^k can then be written as

$$d(P_j^k, P_i^k) = \sqrt{\sum_{n=1}^{p^D} (P_{j_n}^k - P_{i_n}^k)^2}, \quad (5)$$

where $P_{j_n}^k$ is the n th pixel intensity in patch P_j^k . Such a metric can easily be plugged in the tonal kernel of the BF, instead of the squared difference of pixel intensities. This leads to the update

$$\hat{u}_i^{k+1} = \frac{\sum_{j \in N_i} w_{ij} \Psi' \left(d^2(P_j^k, P_i^k) \right) \hat{u}_j^k}{\sum_{j \in N_i} w_{ij} \Psi' \left(d^2(P_j^k, P_i^k) \right)}. \quad (6)$$

D. Variance Stabilizing Transformation in the Patch-Based Bilateral Filter

The presence of a tonal kernel in the above-mentioned filters makes them optimal only for additive Gaussian noise with uniform variance. In the case of Poissonian data, the noise variance of each pixel depends on its intensity. This would require a local adaptation of the tonal kernel width. In order to optimally filter Poissonian noise, we propose to use Anscombe's and Fisz's VSTs. Using Anscombe's VST to filter an image with Poissonian noise consists of three easy steps: i) transform the data before filtering, ii) filter the data as if it was Gaussian, and iii) apply the inverse transformation.

The case of Fisz's VST is slightly more complex, as it involves sums and differences of Poisson variables. This particular form can however be elegantly exploited in order to integrate the transformation directly into the tonal kernel of the filters. The update then becomes

$$\hat{u}_i^{k+1} = \frac{\sum_{j \in N_i} w_{ij} \Psi' \left(\zeta^2(\hat{u}_i^k, \hat{u}_j^k) \right) \hat{u}_j^k}{\sum_{j \in N_i} w_{ij} \Psi \left(\zeta^2(\hat{u}_i^k, \hat{u}_j^k) \right)}. \quad (7)$$

It is noteworthy that constant term $\zeta(\lambda_1, \lambda_2)$ in (2) is dropped for the obvious reason that the argument of the tonal kernel should have a Gaussian distribution but not necessarily a zero mean. We can indeed distinguish two cases. If $\lambda_1 = \lambda_2$ is true, then $\zeta(\lambda_1, \lambda_2) = 0$ and $\zeta(X_1, X_2) \rightarrow N(0, 1)$. On the other hand, if λ_1 and λ_2 sufficiently differ from each other, then it is highly probable that $\zeta(X_1, X_2)$ will fall in the tails of $N(0, 1)$. This leads the tonal kernel to behave exactly as expected: its value should be high when considering two pixels having similar intensities and low if the pixels in either case, the knowledge of λ_1 and λ_2 is not necessary.

As to the PBBF, the update rule (6) remains unchanged, except that the distance between two patches is redefined as

$$d(P_j^k, P_i^k) = \sqrt{\sum_{n=1}^{p^D} (\zeta(P_{j_n}^k, P_{i_n}^k))^2}. \quad (8)$$

III. EXPERIMENTS & RESULTS

A. Benchmark Image Denoising

The task of assessing the filtering performances on images acquired on medical devices is difficult because the ground truth, the noise free image, remains unknown. For this reason, the denoising performances of each combination of VST and algorithm have been tested on a 64×64 benchmark gray level image. This image is a pattern (see Fig. 1, left) consisting in a combination of constant plateaus, edges, and gradients. Together, these regions summarize all types of types of difficulties encountered in medical images.

This pattern is slightly blurred in order to avoid sharp edges are never observed in real world images.

The filtering performances are measured by the RMSE:

$$RMSE = \sqrt{\frac{1}{M \sum_{i=1}^N v_i} \sum_{m=1}^M \sum_{i=1}^N v_i (\hat{u}_i^k - u_i)^2}, \quad (9)$$




Fig. 1. From left to right: Blurred benchmark image without noise, plateaus mask, edges mask, ramps mask.

where M is the number of trials and m then trial index. The weights v_i either have a value of $\{0, 1\}$ and are used to define the masks. These masks allow to evaluate the RMSE over the selected areas of the image (whole image, plateaus or ramps (Fig. 1)). More details about the benchmark image, its generation and the performances evaluation can be found in [17].

The PBBF and BF have been used with Anscombe's and Fisz's VSTs to denoise the benchmark image with added poisson noise. The denoising performances of each algorithm are then evaluated and compared. First, the parameters of the PBBF and BF are optimized for the RMSE on a training image. Then, these parameters are used to filter the same benchmark image, with additional noise generated with the same noise model. This process is repeated over 100 images. The results presented in Table I are the mean and standard deviation of the RMSE between the original image and the denoised test image for the PBBF and the BF with both VST. They should be compared with the results of the same algorithm applied on the image with poisson noise, without any VST (gaussian filter). On this table, L controls the length of the square patch in one direction from the central pixel: $D = 2L + 1$, and n is the optimal number of iterations. These results show that the two VSTs improve the results over the same algorithm for gaussian noise. Independently of the chosen VST, and even without any VST, the patch-based algorithms outperform the BF on all masks. Even if the RMSE of the gaussian PBBF and the VSTs PBBF denoised images are close, a visual examination of those shows that using the VSTs leads to a improvement: the ramps and constant areas have a lower intensity in the case of the gaussian PBBF.

B. PET Image Denoising

In order to complete its performance assessment, the PBBF has been applied to a real PET image. As the ground truth image is not available, the RMSE cannot be computed and the parameters have been adjusted by visual inspection of the residuals $r_i = \hat{u}_i - f_i$. Optimal parameter values lead to a residual image that contains mainly the noise component and few or no remaining structural elements. If the filter alters the underlying signal of the image, parts of it, such as edges, will appear in the residual image. Figure 3 shows the results of the experiment. The result of the PBBF in panel

(b) seems to perform well and smoothes the constant areas of the image while preserving the structure. In contrast, panel (d) shows that the BF smoothes the edges, as confirmed by visual inspection of the residuals in (e). The residuals of the PBBF that are illustrated in (c) contain less structural elements.

IV. CONCLUSION

This paper shows how Anscombe's and Fisz's variance stabilizing transformations can be cast within the framework of the patch-based bilateral filter. It is shown that with these transformations, the patch-based bilateral filter efficiently succeeds in attenuating Poissonian noise in images. Experiments show that the patch-based bilateral filter clearly outperforms the classical bilateral filter both qualitatively and quantitatively.

REFERENCES

- [1] R. Charnigo, J. Sun, and R. Muzic. A semi-local paradigm for wavelet denoising. *IEEE Trans Image Process*, 15(3):666–677, Mar 2006.
- [2] J. W. Lin, A. F. Laine, and S. R. Bergmann. Improving pet-based physiological quantification through methods of wavelet denoising. *IEEE Trans Biomed Eng*, 48(2):202–212, Feb 2001.
- [3] V. Bruni, B. Piccoliand, and D. Vitulano. Wavelets and partial differential equations for image denoising. *Electronic Letters on Computer Vision and Image Analysis*, 6:36–53, 2008.
- [4] S. Osher, M. Burger, Goldfarb D., Xu J., and W. Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.*, 4:460–489, 2005.
- [5] J.M. Sanches, Nascimento J.C., and J.S. Marques. An unified framework for bayesian denoising for several medical and biological imaging modalities. *Conf Proc IEEE Eng Med Biol Soc*, 1:6267–6270, 2007.
- [6] H. Takeda, S Farsiu, and P Milanfar. Kernel regression for image processing and reconstruction. *IEEE Transactions on Image Processing*, 16:349–365, 2007.
- [7] R.-H. Huang and X.-H. Wang. Image denoising via gradient approximation by upwind scheme. *Signal Processing*, 88:69–74, 2008.
- [8] P. Mrazek, J. Weickert, and Bruhn A. *Geometric Properties for Incomplete data*, chapter On Robust Estimation and smoothing with Spatial and Tonal Kernels, pages 335–352. 2006.
- [9] C.K. Chu, I. Glad, F. Godtliebsen, and Marron J.S. Edge-preserving smoothers for image processing. *Journal of the American Statistical Association*, 93,no.442:526556, 1998.
- [10] M. Elad. On the origin of the bilateral filter and ways to improve it. *IEEE Trans. Image Processing*, 11:11411151, 2002.
- [11] C. Tomasi and R. Manduchi. Bilateral ltering for gray and color images. In *6th Int. Conf. Computer Vision*, pp.839846, 1998.
- [12] A. de Decker, J.A. Lee, and M. Verleysen. Performance assessment of patch-based bilateral denoising. In *MIAD2009, First International Workshop on Medical Image Analysis and Description for Diagnosis Systems*, in *BIOSTEC2009, International Joint Conference on Biomedical Engineering Systems and Technologies*, Porto, pp. 52-61, 2009.

TABLE I
RMSE AND STD OF THE RMSE FOR EACH MASKS OVER 100 BENCHMARK IMAGES

	image		plateaus		edges		ramps		σ	ρ	L	n
	RMSE	std	RMSE	std	RMSE	std	RMSE	std				
Noisy image	5.7304	0.0897	6.2386	0.1183	7.2793	0.2354	4.9782	0.1302				
PBBF Fisz	2.6228	0.0870	3.5918	0.1358	3.9521	0.2806	1.2779	0.0988	0.7905	4.4577	1	2
PBBF Anscombe	2.664	0.0879	3.6474	0.1369	4.0078	0.2774	1.3338	0.0962	1.1457	4.1765	1	2
PBBF Gaussian	2.8906	0.1103	3.9499	0.1519	4.6985	0.2941	1.4041	0.1569	5.1831	6.0464	1	4
BF Fisz	3.7066	0.0804	4.8138	0.1171	5.0725	0.2419	2.4098	0.1012	1.3381	1.2516	0	10
BF Anscombe	3.693	0.0798	5.0471	0.1256	4.7222	0.2349	2.2154	0.1019	1.8432	1.2558	0	10
BF Gaussian	3.8701	0.08501	5.1491	0.1354	4.7355	0.2290	2.5687	0.1162	16.7582	1.3726	0	4




Fig. 2. Example of denoised images: from left to right. (a) Original image. (b) BF denoising with Fisz transform: ($RMSE = 3.7066$). (c) BF denoising with Anscombe transform: ($RMSE = 3.693$). (d) Gaussian BF denoising: ($RMSE = 3.8701$). (e) Noisy image: ($RMSE = 5.7304$). (f) PBBF denoising with Fisz transform: ($RMSE = 2.6222$). (g) PBBF denoising with Anscombe transform: ($RMSE = 2.664$). (h) Gaussian PBBF denoising: ($RMSE = 2.8906$).




Fig. 3. PET phantom image: from left to right. (a) Original PET image. (b) PBBF denoised image. (c) Residuals of the PBBF image. (d) BF denoised image. (e) Residuals of the BF image.

- [13] C. Kervrann and A. Trubuil. An adaptive window approach for poisson noise reduction and structure preserving in confocal microscopy. In A. Trubuil, editor, *Proc. IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 788–791 Vol. 1, 2004.
- [14] C. Kervrann and J. Boulanger. Optimal spatial adaptation for patch-based image denoising. *15(10):2866–2878*, 2006.
- [15] P. Fryzlewicz and G. P. Nason. A haar-fisz algorithm for poisson intensity estimation. *Journal of Computational and Graphical statistics*, 13, Number 3:621–638, 2004.
- [16] B. Zhang, B. Zhang, M.J. Fadili, and J.-L. Starck. Multi-scale variance stabilizing transform for multi-dimensional poisson count image denoising. In M.J. Fadili, editor, *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP 2006*, volume 2, pages II–II, 2006.
- [17] J.A. Lee, X. Geets, V. Gregoire, and A. Bol. Edge-preserving filtering of images with low photon counts. *IEEE Transactions on pattern analysis and machine intelligence*, 30:1014–1027, 2008.