Adaptive Aggregation of Multiple Denoisers with Mismatched Correlated Noise Models

Andrea Corsini and Alessandro Foi
Signal Processing Research Centre, *Tampere University*, Finland
firstname.lastname@tuni.fi

***Abstract*—Denoisers for stationary correlated noise require explicit or implicit knowledge of the noise power spectral density (PSD). However, the PSD characterizes the noise correlation only as long as this is stationary. Yet, even when images are corrupted by noise with nonstationary correlation, denoisers can still assume a specific PSD, and successfully filter locally on regions where the noise correlation matches the model provided by the PSD. Elsewhere, the denoising is ineffective and it can introduce distortions and artifacts. One may then want to combine multiple denoised images obtained for a bank of PSDs spanning the diverse correlations found locally over the image, such that each pixel of the final output is obtained using a relevant PSD. However, when the nonstationary correlation is signal-dependent—as in many cases of practical relevance—this combination is not trivial, because the signal is unknown, and so are the regions with matched noise model. We introduce an effective aggregation method specifically designed for signal-dependent correlation, where locally adaptive weights reward the statistical compatibility of estimates and corresponding residual noise with the assumed model. Our adaptive aggregated estimate is close visually and in terms of PSNR to that with oracle weights based on the noise-free image.**

***Index Terms*—Image denoising, nonstationary, correlated noise, adaptive aggregation.**

I. Introduction

Image denoising is often performed assuming that the noise is globally stationary. When this assumption fails, one has to model the nonstationarity, like with heteroskedastic noise, whose variance can be signal-dependent (e.g., shot noise) or spatially variant (e.g., due to photoresponse nonuniformity). Spatial correlation too may vary, like in X-ray images where electronic (background), structure, and quantum (foreground) noises have different normalized spectra [1, 2], and in analog photographic film due to the different overlap of sensitized grains in bright vs. dark regions [3, 4]. Nonstationary correlation can even be caused by simple data transformations such as clipping [5] and variance-stabilizing transformations (VSTs) [6, 7]. VSTs are commonly adopted to transform data with signal-dependent noise variance into data with homoskedastic noise; being nonlinear, VSTs can distort the noise correlation and result in signal-dependent covariances [8].

Denoising of nonstationary correlated noise is relatively understudied; to the best of our knowledge, there is no image denoiser specifically designed for nonstationary correlated noise. Denoisers designed for stationary correlated noise [9, 10] are instead relatively well established, and they typically require knowledge of the noise power spectral density (PSD) to effectively discriminate between signal and noise over different frequencies. However, whenever the noise is nonstationary, the PSD definition as variance of the Fourier-transformed noise does not characterize the noise correlation (Section II). Therefore, denoisers for stationary correlated noise are formally inappropriate for dealing with nonstationary correlation. In practice, they can still perform filtering given a set of specific PSDs used separately assuming local stationarity [11, 12]. They successfully filter the noise of regions whose PSD matches their assumed PSD, but they leave significant portions of noise and artifacts in regions with mismatched noise PSD (Section III).

Unlike previous works that considered the filtering of noise with covariance being locally stationary with respect to the coordinate [13], we focus on noise that is locally stationary with respect to the clean signal intensity (signal-dependent), as it is a relevant model to all the aforementioned scenarios. To this end, we propose an adaptive procedure that identifies relevant regions within a given denoising estimate, and aggregates a limited number of denoising estimates, resulting in an efficient method for the denoising of signal-dependent nonstationary correlated noise (Section IV). The proposed method achieves better denoising than any single estimate when taken separately, but especially its performance is close visually and by PSNR to that of an oracle that aggregates estimates based on the knowledge of the true image (Section V).

II. Locally stationary correlated noise

We consider the following observation model

$$z(x) = y(x) + \sigma(y(x))\eta(x), \quad \eta(x) \sim \mathcal{N}(0,1), \qquad (1)$$

where $x \in \Omega \subset \mathbb{Z}^2$ is the image coordinate, $y : \Omega \to Y \subseteq \mathbb{R}$ is an unknown deterministic clean image, $\sigma : Y \to \mathbb{R}^+$ is a deterministic univariate function that determines the standard deviation of the noise in z, and η is a nonstationary spatially correlated standard Gaussian noise field. We model[1] η as the filtering of an uncorrelated standard Gaussian white noise $\nu(x) \sim \mathcal{N}(0,1)$ with a signal-dependent kernel $g_{y(x)}$ normalized such that $\|g_{y(x)}\|_2 = 1$ for every $y(x)$:

$$\eta(x) = \int_{-\infty}^{+\infty} \nu(\xi)\, g_{y(x)}(x-\xi)\mathrm{d}\xi. \qquad (2)$$

This work was supported by the Research Council of Finland (336357, 351623).

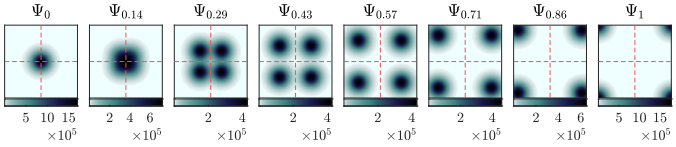

Fig. 1. Slices of an extreme example of Ψ_y (3). The PSD function is more and more high-frequency as y increases, as in the formula $\Psi_y(u) = A_y e^{\frac{u_1^2+u_2^2+(y|\Omega|)^2/2}{800}} \cosh \frac{u_1 y|\Omega|}{800} \cosh \frac{u_2 y|\Omega|}{800}$, where $u \in \Omega$ is the frequency coordinate, A_y is factor such that $\|\Psi_y\|_1 = |\Omega|^2$.

Therefore, η is not (wide-sense) stationary, having constant variance $\mathrm{var}\{\eta(x)\} = 1$ but shift-variant covariance $\mathrm{cov}\,[\eta(x_1), \eta(x_2)] \neq \mathrm{cov}\,[\eta(x_1+\xi), \eta(x_2+\xi)]$, $x_1, x_2, \xi \in \Omega$. This covariance cannot be represented as a function only of the lag parameter ξ, and no PSD can characterize the noise correlation.

As the autocovariance function of η is signal-dependent, it can be considered locally shift-invariant for a given value of y and expressed as a function $R_y : \Omega \to [-1, 1]$ of ξ. This allows one to define a local PSD as a function[2] given y as

$$\Psi_y := \mathrm{var}\{\mathcal{F}[\eta]|y\} = \mathcal{F}[R_y]\,|\Omega| = |\mathcal{F}[g_y]|^2|\Omega|, \qquad (3)$$

where $\mathcal{F}$ is the Fourier transform over Ω, and $|\Omega|$ is the domain cardinality. Consequently, for every $x \in \Omega$, η can be expressed as $\eta(\zeta) \approx \left[\mathcal{F}^{-1}\left[\mathcal{F}[\nu]\,\sqrt{\Psi_{y(x)}|\Omega|^{-1}}\right]\right](\zeta)$ locally for ζ in a neighborhood of x; this approximation is exact for $\zeta = x$.

III. Denoisers for stationary correlated noise applied to locally stationary noise

We analyze the performance of denoisers designed for stationary correlated noise on images z following the nonstationary noise model (1)-(2). We denote these denoisers as an operator D that produces an estimate $\hat{y} = \mathrm{D}(z, \mathcal{P})$ of y from z given an input noise PSD $\mathcal{P}$.

For a more direct illustration, we consider homoskedastic noise with constant $\sigma(y) \equiv \sigma \in \mathbb{R}^+$ and a particularly challenging signal-dependent covariance Ψ_y shown in Fig. 1. Fig. 2 shows an example of y, the corresponding z, as well as two estimates $\hat{y}$ obtained with the BM3D denoiser for stationary spatially correlated noise [9] with input PSD $\mathcal{P} = \sigma^2\Psi_{\check{y}}$, with $\check{y}$ equal either to 0.26 or 0.66. The denoising quality is satisfactory in regions where y is approximately equal to the $\check{y}$ value of the assumed PSD $\sigma^2\Psi_{\check{y}}$ used as input for D. Elsewhere, there are artifacts, oversmoothing of image details, and leftover noise.

We confirm this by considering the weighted PSNR

$$\mathrm{WPSNR}(\hat{y}, y, \omega_{\check{y}}) = 10 \log_{10}\left(\frac{\max_{x\in\Omega} \omega_{\check{y}}(x) y(x)}{\frac{\sum_{x\in\Omega} \omega_{\check{y}}(x)(\hat{y}(x)-y(x))^2}{\sum_{x\in\Omega} \omega_{\check{y}}(x)}}\right), \qquad (4)$$

[1]The model (1)-(2) is general enough to sufficiently cover several cases, like those described in Section I. More sophisticated models may be considered based on the specific application.

[2]The functional dependency on the 2-D frequency coordinate is implicit.

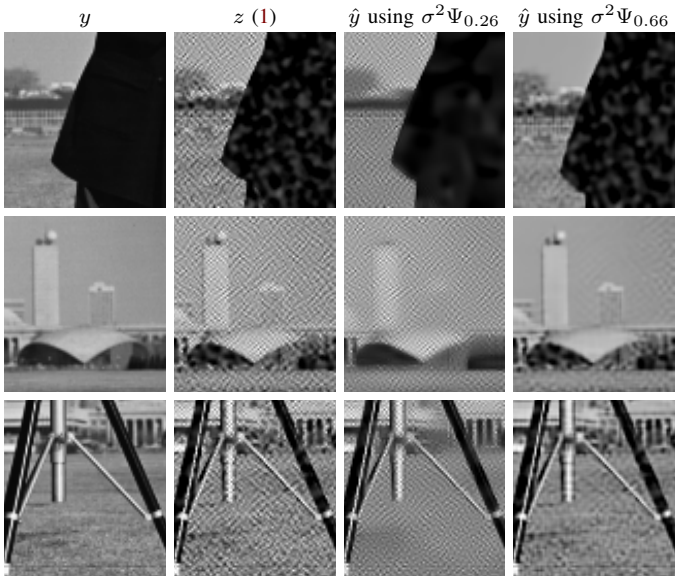

Fig. 2. Left to right, details of: *Cameraman* y, noisy image z (1) (PSNR 20.11 dB) with $\sigma = 0.1$ and Ψ_y from Fig. 1, and two estimates $\hat{y}$ (PSNR 21.83 dB, 24.75 dB) obtained using input noise PSD $\sigma^2\Psi_{\check{y}}$ with $\check{y} = 0.26, 0.66$. In both estimates, the denoising is satisfactory where $y \approx \check{y}$, featuring artifacts and leftover noise elsewhere.

where $\omega_{\check{y}} : \Omega \to [0, 1]$ denotes the metric weight defined as

$$\omega_{\check{y}}(x) = e^{-\frac{1}{2}\left(\frac{y(x)-\check{y}}{\varepsilon}\right)^2}. \qquad (5)$$

Essentially, $\omega_{\check{y}}$ favors pixels in regions such that $y(x) \approx \check{y}$. We evaluate (4) on estimates $\hat{y}$ obtained for input PSD $\Psi_{\check{y}}$ sweeping all combinations of $\check{y}$ and $\check{y}$ with $0 \leq \check{y}, \check{y} \leq 1$. Fig. 3a reports these $\mathrm{WPSNR}(\hat{y}, y, \omega_{\check{y}})$ normalized between 0 and 1 for every value of $\check{y}$, $\varepsilon = 0.012$. Rather than the global maximum of WPSNR, we analyze which $\check{y}$ gives the maximum score for each $\hat{y}$, i.e. $\check{y}^\diamond(\check{y}) = \arg\max_{\check{y}} \{\mathrm{WPSNR}(\hat{y}, y, \omega_{\check{y}})|\Psi_{\check{y}}\}$. We can see in Fig. 3a that indeed $\check{y}^\diamond(\check{y})$ is distributed along the diagonal $\check{y} = \check{y}$, i.e. each $\hat{y}$ is best locally where $y \approx \check{y}$. Similar considerations are found to be general for other images and nonstationary covariances, and also valid for the BLS-GSM denoiser [10].

Despite localized good quality, none of these estimates is globally satisfactory, due to the presence of regions having mismatched noise model. Some estimates can have better global (unweighted) PSNR than others, as one can see in Fig. 3b. However, even the best one suffers from the artifacts and leftover noise seen in Fig. 2.

IV. Proposed method

We propose a denoising procedure for images z (1) that aggregates a (typically small) set of N estimates $\{\hat{y}_i\}_{i=1}^N$ based on adaptive weight maps $w_i^\star$ as in

$$\hat{y}^\star(x) = \frac{\sum_{i=1}^N w_i^\star(x)\hat{y}_i(x)}{\sum_{i=1}^N w_i^\star(x)}, \qquad (6)$$

where $\hat{y}_i = \mathrm{D}(z, \sigma^2(\check{y}_i)\Psi_{\check{y}_i})$ with $\check{y}_i = z_{;\frac{1}{N}(i-\frac{1}{2})}$, $z_{;\beta}$ being the β-quantile of z. We expect this to select $\check{y}_i$ such that $\sigma^2(\check{y}_i)\Psi_{\check{y}_i}$ characterizes the local noise PSD of substantial portions of z.

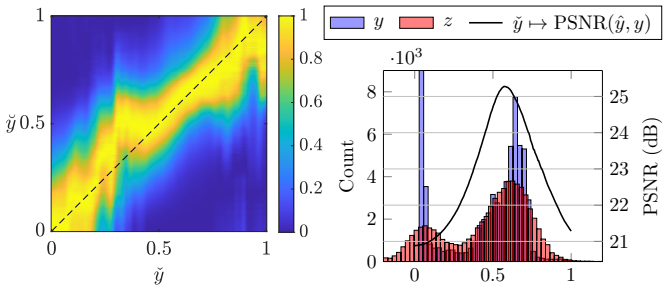

(a) Normalized $\mathrm{WPSNR}(\hat{y}, y, \omega_{\check{y}})$ (4)

(b) Global $\mathrm{PSNR}(\hat{y}, y)$

Fig. 3. Metrics for estimates $\hat{y}$ obtained denoising z (Fig. 2) as described in Section III. Left: $\mathrm{WPSNR}(\hat{y}, y, \omega_{\check{y}})$ (4) as a function of $\check{y}$, which determines the input $\Psi_{\check{y}}$ for each $\hat{y}$, and $\check{y}$, which determines the weight $\omega_{\check{y}}$ (5), and in turn, the region $y \approx \check{y}$, normalized for each $\check{y}$. High scores are distributed around the diagonal, that is, the denoising is more satisfactory in regions $y \approx \check{y} \approx \check{y}$, as the local PSD of such regions matches $\Psi_{\check{y}}$. Right: global PSNR between $\hat{y}$ and y as a function of $\check{y}$. Highest scores are obtained for estimates given by $\check{y} \approx 0.6$, i.e. near the rightmost peaks of the histograms (overlaid in the plot) of y and z. Such value corresponds approximately to the mean of the bright sky in *Cameraman*. As we can see from y and Ψ_y in Fig. 1 and Fig. 2, the higher PSNRs are due to such area having content and noise strong in low and high frequencies, respectively. The PSNRs at the leftmost histogram peaks are lower, as they correspond to an area with both noise and content (the dark coat in *Cameraman*) strong in the low frequencies, resulting in a harder case to denoise than the sky area.

From the analysis in Section III, we see that the analyzed denoiser performances are locally satisfactory in regions such that $y(x) \approx \check{y}_i$. Hence, we can construct $w_i^\star$ similarly to ω (5), by rewarding estimate pixels such that $\hat{y}_i(x) \approx \check{y}_i$ as in

$$w_i'(x) = \phi_{\mathrm{a}}\left(\hat{y}_i(x); \check{y}_i, \tfrac{\check{y}_i - \check{y}_{\max\{i-1,1\}}}{3}, \tfrac{\check{y}_{\min\{i+1,N\}} - \check{y}_i}{3}\right), \quad (7)$$

where $\phi_{\mathrm{a}}(\cdot; \mu, \varsigma_l, \varsigma_r)$ is the the asymmetric[3] Gaussian probability density function [14] with mean μ, left and right standard deviations ς_l and ς_r. Notice that for $\varsigma_l \to 0$ and $\varsigma_r \to 0$ ϕ_a approaches a Dirac impulse δ_μ at μ. Hence, $w_i'(x) \xrightarrow[N\to\infty]{} \delta_{\check{y}_i}(\hat{y}_i(x))$, and the aggregation $\{w_i'(x)\}_i^N$ reduces for each x to an estimate selection. Nevertheless, similar results to such selection can be efficiently reached by convex aggregation of few (small N) estimates [15–17]. However, a limitation of w_i' (7) is that it rewards also estimates having $\hat{y}_i(x) \approx \check{y}_i$ while $\hat{y}_i(x) \not\approx y(x) \not\approx \check{y}_i(x)$ due to significant denoising artifact introduced by a mismatched noise model (Fig. 2).

To overcome this limitation, we observe that, if a denoiser is locally exact giving an estimate $\hat{y}(x) = y$ for x in some neighborhood U and for some value of y, then the local PSD of the residual noise $z - \hat{y}$ should match $\sigma^2(y)\Psi_y$ in U. We use a bank of filters $\{h_k\}_{k=1}^K$ as a gauge for the match between the residual noise and the given $\sigma^2(\hat{y})\Psi_{\hat{y}}$. Specifically, for $\|h_k\|_2 = 1$, the variances of the filter output should satisfy

$$v_k(y) := \mathrm{var}\{z * h_k | y\} = \sigma^2(y) \left\| |\mathcal{F}[h_k]|^2 \Psi_y \right\|_1 |\Omega|^{-2}. \quad (8)$$

Using multiple kernels h_k helps to disambiguate weak responses of poorly filtered noise from those caused by non-overlapping frequency bands of h_k and η. As we have only a single realization of z, variances are computed as sample estimates over small windows. Notice that, as models are signal-dependent, we also need an estimate of the mean which shall be used in place of y in (8). Hence, for each of the N denoised estimates $\hat{y}_i$, we form K mean-variance pairs $\{m_i(x), s_{i,k}^2(x)\}$, where

$$m_i(x) = \frac{1}{L^2} \sum_{\xi \in W_x} (\hat{y}_i * l)(\xi),$$
$$s_{i,k}^2(x) = \mathrm{sample\ var}_{\xi \in W_x} \{[(z - \hat{y}_i) * h_k](\xi)\},$$

W_x denotes an $L \times L$ sliding window referenced at x, and l is a low-pass filter adopted to mitigate leftover noise and artifacts from mismatched noise models. We require $\sum_x l(x) = 1$. We assign a high or low score to each estimate pair $\{m_i(x), s_{i,k}^2(x)\}$ whose sample variance $s_{i,k}^2(x)$ respectively matches or mismatches the variance $v_k(m_i(x))$ given by (8) based on a scoring function ϑ_{v_k} defined as an asymmetric Gaussian

$$\vartheta_v(m, s^2) = \phi_{\mathrm{a}}\left(s^2; v(m), \tfrac{v(m)}{3}\left[1 - \tfrac{\chi^2_{n;\frac{\alpha}{2}}}{n}\right], \tfrac{v(m)}{3}\left[1 - \tfrac{\chi^2_{n;1-\frac{\alpha}{2}}}{n}\right]\right), \quad (9)$$

from which we conveniently omit the subscripts i, k, and the coordinate x. Here, the left and right standard deviations of ϕ_{a} depend on the $1-\alpha$ confidence interval of the χ_n^2 distribution with $n = L^2 - 1$ degrees of freedom, since this can approximate[4] the distribution of $\frac{n s_{i,k}^2(x)}{v_k(m_i(x))}$. Then, denoting by $\mathcal{X}_{W_x} : \Omega \to \{0,1\}$ the indicator function of W_x, we construct the weight map $w_i^\star$ used in (6) by averaging the product of pair scores across the K filters from overlapping windows as

$$w_i'' = \frac{\sum_x \mathcal{X}_{W_x} \prod_{k=1}^K \vartheta_{v_k}\left(m_i(x), s_{i,k}^2(x)\right)}{\sum_x \mathcal{X}_{W_x}}. \quad (10)$$

Finally, we propose the joint weighting strategy where $w_i^\star$ in (6) is the product of w_i' (7) and w_i'' (10), as each weight factor may individually identify mismatched regions that the other misses, as we shall see in Section V and Fig. 4.

V. Experiments

We consider the noisy image z in Fig. 2 and evaluate the proposed adaptive aggregation (6) in three alternative forms: a *Basic′* solution that uses exclusively w_i' (7), a *Basic″* solution that uses w_i'' (10), and the *Full* solution consisting of the product of w_i' (7) and w_i'' (10). We also consider an *Oracle* solution, where the weights w_i' (7) have y instead of $\hat{y}_i$.

We fix the sliding window size to $L = 7$ and $K = 2$ filters for the variance. Specifically, a high-pass kernel $h_1 = H_1 \psi \psi^\top$ and a low-pass kernel $h_2 = H_2 \varphi \varphi^\top$, where ψ and φ are respectively the 1-D Daubechies wavelet and scaling functions with three vanishing moments, and H_k such that $\|h_k\|_2 = 1$. We choose the filter for the mean as $l = h_2 / \sum h_2$.

The results in Fig. 4 confirm that aggregating few $\{\hat{y}_i\}_i$ performs better than any single estimate $\hat{y}_i$ (cf. Fig. 3b). However, the *Basic′* solution still has artifacts due to the

[3] The asymmetry follows from $\{\check{y}_i\}_{i=1}^N$ possibly being nonequispaced.

[4] This is only a coarse approximation primarily because the samples on which $s_{i,k}^2$ is calculated are not independent due to the noise being correlated.

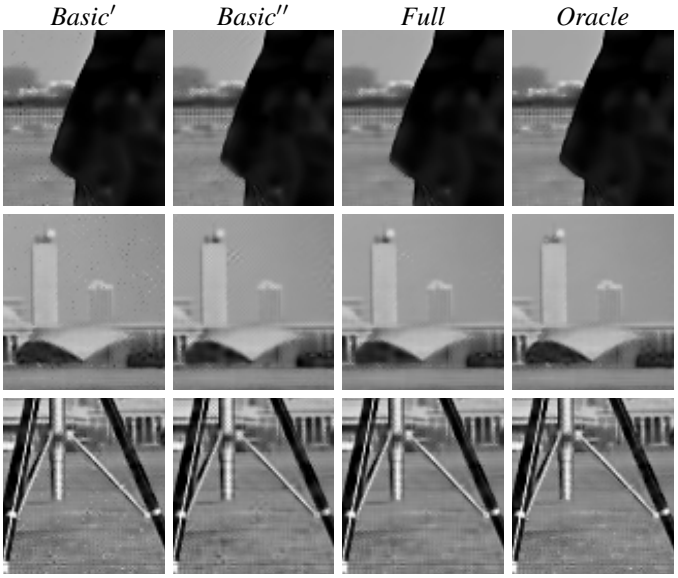

Fig. 4. Details from the alternative aggregations (Section V) used on $N=5$ estimates $\{\hat{y}_i\}_i$ given z from Fig. 2. The *Full* aggregation (PSNR 28.53 dB) is free from artifacts visible in *Basic′* (PSNR 28.30 dB) and *Basic″* (PSNR 28.34 dB), and it is visually similar to *Oracle* (PSNR 28.91 dB).

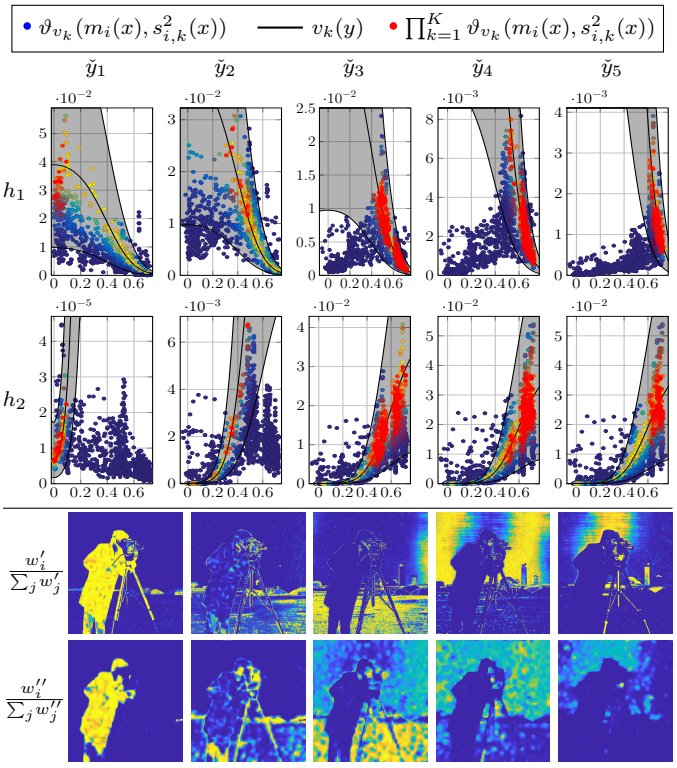

Fig. 5. Estimate pair scores and the weight maps w'_i (7) and w'' (10), used for aggregating the estimates $\hat{y}_i$ that lead to the results shown in Fig. 4. Top: scatterplots showing the sample pairs $\{(m_i(x), s^2_{i,k}(x))\}_{i,k}$ from several sliding windows W_x, and their assigned individual score $\vartheta_{v_k}(m_i(x), s^2_{i,k}(x))$ (9) (high score yellow, low score blue). Pairs whose $s^2_{i,k}$ is closer to the model $v_k(m_i)$ (8) get a higher score. The opacity of the red alpha-blended scatterplot shows the combined scores as the product of individual scores coming from h_1 and h_2. The combined score helps to disambiguate sample pairs that score high only in one of the two models $v_k(y)$ due to the noise frequencies on $\hat{y}_i$ that vanishes outside the bandwidth of the filter. Bottom: weights w'_i (7) and w''_i (10). For visual clarity, scatterplots show an independent subsample (2%) of the total sample pairs (62 500) that were used to compute the weights w''_i.

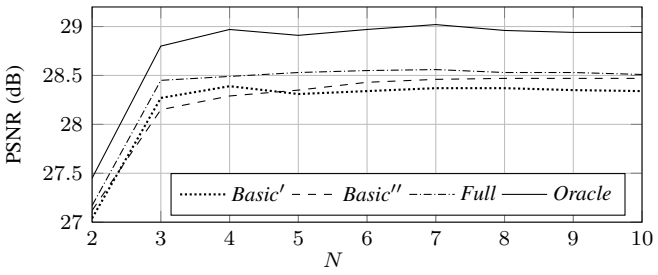

Fig. 6. PSNR of the *Basic′*, *Basic″*, *Full* and *Oracle* solutions for increasing number N of aggregated estimates. The PSDs used in the denoising are chosen automatically based on the histogram of z (see Section IV). For $N=1$, all solutions reduce to $\hat{y}_1$ given by $\hat{y}_1 = \text{median}_x\{z(x)\}$ (PSNR 25.10 dB).

residual noise in $\hat{y}_i$ that confuses the test against $\check{y}_i$ in w'_i (7). In fact, the *Oracle* solution is free from such artifacts, as the clean y is instead tested against $\check{y}_i$. *Basic″* is also free from the artifacts found in *Basic′*, as w''_i (10) penalizes regions of $\hat{y}_i$ whose residual noise has estimated variances not matching the assumed ones $\{v_k\}_k$ (8). Yet, different artifacts appear due to variations in the estimated mean m used to select the variance $v_k(m)$ to be tested. However, *Full* is free from most of the artifacts of *Basic′* and *Basic″*, and it approaches the visual quality of the *Oracle* solution.

Fig. 5 visualizes the effects of using multiple filters $\{h_k\}$ in the match of the variances, which lead to w''_i (10) used in *Basic″* and *Full* aggregations. In particular, we show the score assigned to estimated mean-variance pairs $\{(m_i(x), s^2_{i,k}(x))\}$ by ϑ_{v_k} (9) for several windows. We can see that without multiple filters h_k, we would favor many more pairs, including those from windows with a mismatched noise model which happened to have a high score only because the filter may have had no frequency response for that particular noise spectrum. Fig. 5 also exhibits the different image structures favored by the two weight maps w'_i and w''_i used in Fig. 4.

Fig. 6 shows that already a few estimates are sufficient to attain a substantial PSNR improvement for all aggregation solutions, thanks to the selection of $\check{y}_i$ values based on the quantiles of z.

VI. Conclusion

We proposed an adaptive aggregation method designed for nonstationary noise with signal-dependent covariance. It efficiently combines multiple estimates obtained from denoisers for stationary correlated noise. The aggregation is based on adaptive weights that reward residual-noise statistics that are compatible with the noise model assumed by the denoiser. Future work can focus on improving and further automatizing the algorithm parameter selection. For example the definition of the filter bank $\{h_k\}_k$ used in the matching of the statistics can be adaptive based on the factorization of the PSD function (3).

REFERENCES

[1] A. Mackenzie, N. W. Marshall, A. Hadjipanteli, D. R. Dance, H. Bosmans, and K. C. Young, "Characterisation of noise and sharpness of images from four digital breast tomosynthesis systems for simulation of images for virtual clinical trials," *Physics in Medicine and Biology*, vol. 62, no. 6, pp. 2376–2397, Mar. 2017.

[2] A. M. M. M. Perez, L. A. S. Lopes, R. F. Caron, B. B. Oliveira, and M. E. Poletti, "Performance evaluation of six digital mammography systems," *Radiation Physics and Chemistry*, vol. 218, p. 111635, May 2024.

[3] A. Newson, J. Delon, and B. Galerne, "A stochastic film grain model for resolution-independent rendering," *Computer Graphics Forum*, vol. 36, no. 8, pp. 684–699, Dec. 2017.

[4] K. Zhang, J. Wang, D. Tian, and T. N. Pappas, "Film grain rendering and parameter estimation," *ACM Transactions on Graphics*, vol. 42, no. 4, pp. 1–14, Aug. 2023.

[5] J. Van Vleck and D. Middleton, "The spectrum of clipped noise," *Proceedings of the IEEE*, vol. 54, no. 1, pp. 2–19, 1966.

[6] B. Efron, "Transformation theory: How normal is a family of distributions?" *The Annals of Statistics*, vol. 10, no. 2, pp. 323–339, 1982.

[7] A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian, "Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data," *IEEE Transactions on Image Processing*, vol. 17, no. 10, pp. 1737–1754, Oct. 2008.

[8] A. Corsini and A. Foi, "Spectra of variance-stabilized correlated noise: Modeling and efficient computation," in *2025 IEEE Statistical Signal Processing Workshop*. Edinburgh, United Kingdom: IEEE, Jun. 2025.

[9] Y. Mäkinen, L. Azzari, and A. Foi, "Collaborative filtering of correlated noise: Exact transform-domain variance for improved shrinkage and patch matching," *IEEE Transactions on Image Processing*, vol. 29, pp. 8339–8354, 2020.

[10] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Transactions on Image Processing*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.

[11] D. Tjøstheim, "Spectral generating operators for non-stationary processes," *Advances in Applied Probability*, vol. 8, no. 4, pp. 831–846, Dec. 1976.

[12] W. Martin and P. Flandrin, "Wigner-Ville spectral analysis of nonstationary processes," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 6, pp. 1461–1470, Dec. 1985.

[13] J. Sills and E. Kamen, "Wiener filtering of nonstationary signals based on spectral density functions," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 3. New Orleans, LA, USA: IEEE, 1995, pp. 2521–2526.

[14] T. Kato, S. Omachi, and H. Aso, "Asymmetric Gaussian and its application to pattern recognition," in *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 2002, pp. 405–413.

[15] A. Nemirovski, "Topics in Non-Parametric Statistics," in *Lectures on Probability Theory and Statistics: Ecole d'Eté de Probabilités de Saint-Flour XXVIII - 1998*, ser. Lecture Notes in Mathematics, P. Bernard, Ed. Springer Berlin Heidelberg, 2000, vol. 1738, pp. 85–277.

[16] F. Bunea, A. B. Tsybakov, and M. H. Wegkamp, "Aggregation for Gaussian regression," *The Annals of Statistics*, vol. 35, no. 4, pp. 1674–1697, Aug. 2007.

[17] A. Goldenshluger, "A universal procedure for aggregating estimators," *The Annals of Statistics*, vol. 37, no. 1, pp. 542–568, Feb. 2009.