Information Fusion 101 (2024) 101999

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

Full length article

Binlets: Data fusion-aware denoising enables accurate and unbiased quantification of multichannel signals

Mauro Silberberg, Hernán E. Grecco *

Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física, Buenos Aires, Argentina
CONICET - Universidad de Buenos Aires, Instituto de Física de Buenos Aires (IFIBA), Buenos Aires, Argentina

ARTICLE INFO

Dataset link: https://github.com/maurosilber/binlets

Keywords:
Wavelets
Denoising
Signal processing
Multichannel
Time series
Images

ABSTRACT

As monitoring multiple signals becomes more cost-effective, combining them through a data fusion-aware denoising method can produce a more robust estimation of the underlying process. Here, we present a method based on the Haar wavelet transform that trades off resolution against accuracy based on statistical significance. By taking advantage of correlations between channels, it offers a superior performance compared to denoising each channel separately. It outperforms standard wavelet methods when the magnitude of interest in the data-fusion process involves a non-linear transformation or reduction of a multichannel signal. We demonstrate its efficacy by benchmarking our method against standard wavelet thresholding for synthetic single and multichannel time series, and a multichannel two-dimensional image. The method has a simple interpretation as an adaptive binning of the signal, and neither requires training data nor specialized hardware to run fast. In addition, a reference Python implementation is available on GitHub and PyPI, making it simple to integrate into any analysis pipeline.

1. Introduction

Monitoring multiple observables simultaneously is becoming widespread, as acquiring and storing them is now more cost-effective. When these observables are a partial view of an underlying process, it can be reconstructed through data fusion. An existing challenge is to obtain a robust estimation of the time evolution of this process of interest. As natural signals are continuous, data fusion-aware denoising techniques could be leveraged to derive precise and accurate results.

Denoising, the removal of noise from a signal, is an important area of research in signal processing as it can increase the accuracy in the quantification of observables. In turn, this has an impact on experimental methods. For instance, in microscopy, better denoising algorithms enable lower-light conditions or shorter exposure times, which allows for longer observation times and reduced toxicity for samples in biological experiments. Additionally, since data visualization is an essential task in data analysis, as exemplified by the well-known Anscombe's quartet [1], denoising improves the reliability of data analysis and interpretation.

Several methods have been developed to address denoising, which can be broadly classified as spatial, transform-based, sparse representation, and patch-based methods, among other categories [2–4]. Transform-based methods expand the signal in a basis in which most coefficients are small and can be discarded. Among these, the wavelet transform has been extensively studied [5] as it is particularly effective for signals with multiscale characteristics. Moreover, it can be performed efficiently using a recursive algorithm that decomposes the signal into approximation and detail coefficients. These detail coefficients correspond to different space-frequency scales, where high-frequency coefficients are highly localized in space.

To discard these coefficients, several thresholding schemes have been proposed [5]. The simpler is hard thresholding, which involves setting coefficients smaller than a given threshold to zero while retaining larger coefficients as they are. For normally-distributed noise, thresholding can be interpreted as a statistical test at a given significance [6]. Nevertheless, in many imaging applications, such as microscopy, signals are corrupted by Poisson or Poisson–Gaussian noise [2]. Variance-stabilizing transforms, such as Anscombe's transform [7], are used to transform the data into an approximately Gaussian distribution, which can then be denoised using standard thresholding methods. Other methods consider how noise is distributed among coefficients for a specific wavelet. For instance, the Haar wavelet is particularly

* Corresponding author.
E-mail address: hgrecco@df.uba.ar (H.E. Grecco).

https://doi.org/10.1016/j.inffus.2023.101999
Received 13 April 2023; Received in revised form 31 July 2023; Accepted 28 August 2023
Available online 1 September 2023

suitable for Poisson denoising [8,9], although it produces discontinuous estimates with block-like artifacts.

Wavelet analysis has been extended in many directions. For denoising, redundant or translation-invariant decompositions have been studied, which average the denoising over several translations of the signal, yielding a smoother estimate [10]. In addition, it has been expanded to handle vector-valued or multichannel signals, such as color images, where the coefficients are vectors and are thresholded as a whole [11,12]. While analyzing each channel independently is possible, as was done in [13] for fluorescence lifetime imaging microscopy (FLIM), we have shown that considering all channels simultaneously achieves much better performance [14].

In recent years, machine learning methods, particularly deep learning, have dominated the scene in terms of denoising [15]. These methods are not based on statistical modeling of the processes involved in the generation and acquisition of a signal, but are *trained* to *learn* a transformation of the signal from pairs of inputs and desired output. In general, this training produces *ad-hoc models* that cannot be reused for samples with different properties, and the recommendation is to train them from scratch [16]. This not only involves a computationally expensive process that requires specialized hardware (GPUs), but also the acquisition of large datasets with ground truth images [17]. As they might be costly, difficult or impossible to obtain experimentally, methods that do not require a ground truth have also been studied [18]. Nevertheless, these black-box *models* tend to be biased toward previously seen structures, generate artifacts [19], and have no *a priori* knowledge of when they will fail.

In this article, we present a data fusion-based method to denoise multichannel and multidimensional signals using a translation-invariant Haar wavelet decomposition. Unlike previous wavelet-based methods, which discard detail coefficients based on their values, we perform a comparison on a different transformation of the approximation coefficients from which the detail coefficients are computed. This new method allows us to take advantage of correlations between channels, offering a superior performance specially when a non-linear transformation of the signal is involved. We show that computing the wavelet decomposition in the original space, but deciding to denoise based on the transformed space, results in substantially better precision and accuracy compared to the standard wavelet approach which denoises each channel separately. Additionally, it is grounded in solid statistical modeling, allowing an understanding of its limitations, and offers a simple interpretation as a spatially-adaptive binning of multichannel and multidimensional data.

In the section *Theory* section, we present a short review of the wavelet transform and the connection between coefficient thresholding and statistical testing. In the *Results* section, we first present a single-channel signal to compare this algorithm to the standard wavelet approach. Then, for a dual-channel signal, we compare the joint denoising enabled by binlets against denoising each channel separately with the standard approach. Lastly, we show the joint denoising of another dual-channel signal using as a target transformation a single-channel observable which is derived from them.

2. Theory

The wavelet transform provides a time-frequency decomposition of a signal, half-way to the Fourier transform. In contrast to the latter, there are many orthogonal wavelet basis functions to choose from, which provide multiresolution approximations of the signal. Then, a non-linear denoising can be performed by setting small coefficients to zero, although there are alternative thresholding methods. An extension that is more suitable for denoising is the translation-invariant dyadic transform, which provides an overcomplete decomposition. Additionally, there are extensions for N-dimensional signals and multichannel or vector-valued signals. An in-depth explanation of this can be found in [10].

Our method, binlets, consists of a translation-invariant wavelet transform with the non-normalized Haar wavelet. In general, we apply it to an $x \in \mathbb{R}^n$ signal, from which we intend to compute a derived $f(x) \in \mathbb{R}^m$ signal. We perform the wavelet decomposition on x, but the decision to threshold each coefficient is performed in the $\mathbb{R}^m$ space. For non-linear $f(x)$, there is an additional condition: the coefficients of the previous levels must also have been thresholded. In this last case, it is equivalent to an adaptive binning of the signal; hence, the name *binlets*. A full description in pseudocode is available in Appendix.

In the following subsections, we provide a brief review of the Haar transform. We show that thresholding is equivalent to performing a statistical test, and the denoised signal results in an adaptive binning of the original signal.

2.1. Haar wavelet and multiresolution analysis

The (non-normalized) Haar wavelet is the simplest wavelet, which computes its coefficients as sums and differences from the original signal $\{x_i\}$. Starting with the assignment $a_i^0 = x_i$, for a given level L, we compute the approximation a_i^L and detail d_i^L coefficients as:

$$\begin{cases} a_i^{L+1} &= a_i^L + a_{i+s}^L \\ d_i^{L+1} &= a_i^L - a_{i+s}^L \end{cases} \tag{1}$$

where $s = 2^L$ is the shift at this level. This can be inverted as

$$\begin{cases} a_i^{L-1} &= (a_i^L + d_i^L)/2 \\ a_{i+s}^{L-1} &= (a_i^L - d_i^L)/2. \end{cases} \tag{2}$$

In Fig. 1, we show this decomposition for $l = 0$, where instead of showing the a_i^1 coefficients, we show $a_i^1/2$.

To produce a multiresolution analysis, this process is recursively repeated for the next level L. If we downsample the approximation signal, keeping only the a_{2i}^L coefficients (as shown in Fig. 1), it results in a decomposition in an orthogonal basis. If not, it results in an overcomplete or redundant decomposition, which makes it translation-invariant.

2.2. Denoising and statistical testing

To denoise a signal, we set the detail coefficients d_i to 0 if they are smaller than a given threshold; that is, if x_i and x_{i+1} are close. Then, when we perform the inverse transform (2), we obtain

$$\hat{x}_i = \hat{x}_{i+1} = \frac{a_i}{2} = \frac{x_i + x_{i+1}}{2}. \tag{3}$$

The reconstructed coefficients $\hat{x}_i$ and $\hat{x}_{i+1}$ are the averages of the original coefficients. This results in an adaptive averaging of the signal, a trade-off between spatial resolution and accuracy, which is schematized in Fig. 1.

To select an appropriate threshold, it is important to understand how the noise in the original signal propagates to the detail coefficients.

When the signal has Gaussian noise, $x_i \sim N(\mu_i, \sigma^2)$, the detail coefficients d_i are also Gaussian: $d_i \sim N(\mu_i - \mu_{i+s}, 2^l\sigma^2)$, where l is the level. To threshold this coefficient, we can test whether it is significantly different from 0 with a Z-score or χ^2 test at a significance level α. For this type of noise, it is more convenient to use the normalized Haar wavelet, which has a $\sqrt{2}$ factor in the forward transform (1), as the detail coefficients have the same variance σ^2 at every level. Then, the usual recipe of using a global threshold λ is equivalent to performing a test at a significance level α. For instance, $\lambda = 2\sigma$ for $\alpha \approx 0.05$.

For signals with Poisson noise, the non-normalized Haar wavelet (1) is more appropriate as each approximation coefficient a_i is also Poisson-distributed. However, in this case, it is not appropriate to use a global threshold for all the detail coefficients. Since we can estimate the variance as x_i (or a_i), we can perform a χ^2 test for d_i, approximating the Poisson x_i as a Gaussian with mean and variance x_i. We threshold $d_i = x_i - x_{i+s}$ when

$$\frac{\text{Mean}[d_i]^2}{\text{Variance}[d_i]} = \frac{(x_i - x_{i+s})^2}{x_i + x_{i+s}} = \frac{d_i^2}{a_i} < \lambda \tag{4}$$

with $\lambda = 2$ for a significance $\alpha \approx 0.05$.

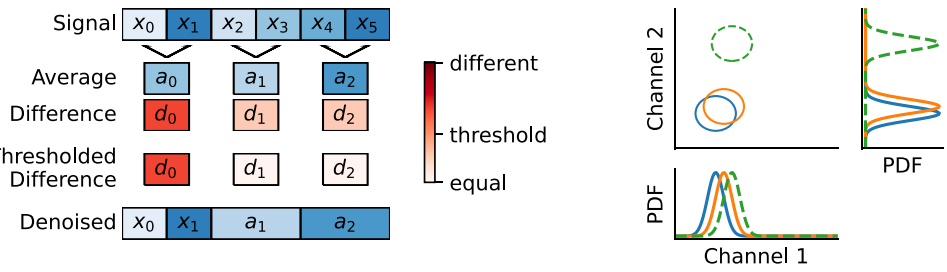

Fig. 1. (left) A schematic one-level decomposition of a signal into average and difference coefficients. Difference coefficients smaller than a given threshold were set to zero. In those cases, the reconstructed signal consists of the average of neighbors at those locations, effectively trading spatial resolution for precision. (right) Three 2D points with their covariance ellipses, and the marginal distributions in each axis. The green point cannot be identified as different from the other two points if only the information from channel 1 is used, as they overlap in that direction.

2.3. Multivariate transform and thresholding

In the case of a multichannel or vector-valued signal $\vec{x}_i$, the same decomposition can be performed on each channel. Then, we obtain vector coefficients $\vec{a}_i$ and $\vec{d}_i$ in (1). Instead of denoising each channel separately, we can leverage the multichannel information to apply a threshold to the whole vector $\vec{d}_i$. As depicted in Fig. 1, the green point is different from the blue and orange ones, when considering their covariance ellipses. Nevertheless, we cannot distinguish between them based on the channel 1 data alone, as shown in the marginal plot below.

In summary, these methods recursively decompose the signal as an average of neighbors. The usual approach is to apply a single global threshold, which works well for Gaussian noise. For a more general noise distribution, if the variance is a function of the mean, a variance-stabilizing transform can be applied to obtain a signal with approximately Gaussian noise. However, they do not leverage information on the correlations between the channels.

3. Results

In this section, we compare binlets against standard wavelet thresholding. First, for a single-channel one-dimensional signal with Poisson noise, where it outperforms the standard approach of using a variance-stabilizing transform. Then, for a multichannel one-dimensional signal with normal noise, to show that thresholding the detail vector jointly is better than considering each channel separately. Finally, for a multichannel signal that is non-linearly transformed to a scalar signal, binlets shows the greatest difference from the standard approach. In this last case, no variance-stabilizing transform is possible, since the variance is not a function of the mean.

To evaluate the performance, we considered the weighted root mean squared error (RMSE) and bias, defined as:

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} Z_i^2} \tag{5}$$

$$\text{bias} = \frac{1}{N}\sum_{i=1}^{N} Z_i \tag{6}$$

where $Z_i = (x_i - x_i^{true})/\sigma_i^{true}$ is the difference between the measured and true values divided by the true standard deviation. These two measures were calculated as a function of the significance level α, which is shown as the equivalent standard deviation σ for a normal distribution.

For the signals, we used a test signal designed in [20], which is widely used in the wavelet literature [10]. Code to reproduce the analysis and figures is available in the GitHub repository: https://github.com/maurosilber/binlets.

3.1. Denoising a single-channel time series

An established way of handling a signal with Poisson noise (Fig. 2) is to use the Anscombe transform [7], which produces a signal whose noise follows approximately a standard normal distribution. Then, as discussed before, we can perform a decomposition with the normalized Haar wavelet and apply a single global threshold to all detail coefficients, which is equivalent to performing a statistical test at a significance level α. Finally, we can apply the inverse Anscombe transform to obtain the denoised signal. Alternatively, we can directly apply the test of Eq. (4) to each detail coefficient using binlets, without using the Anscombe transform.

In Fig. 2, we compared both methods for different mean values of the true signal. The difference between the methods is negligible for high mean values (left) as the Poisson tends towards a normal distribution. In contrast, binlets performs better for lower mean values (right) as we are considering the exact distribution for each coefficient instead of an approximation.

3.2. Denoising a multi-channel time series

Wavelet denoising can easily be extended to multichannel signals. Here, we consider a multichannel or vector-valued signal that alternates between four vectors, as depicted in Fig. 3. For the noise distribution, we used independent standard normal noise, as no (co)variance-stabilization was needed for the standard wavelet approach. In the more general case, a covariance-stabilizing transform might not exist [21]. Hence, the standard approach of applying a single threshold globally would not be statistically optimal. However, we could take into account the covariance between the vector components with binlets.

To denoise this signal, we consider two alternatives: analyzing each channel separately (standard wavelet approach) or thresholding them jointly (binlets). Some of the true vectors (black dots) differ only slightly in one channel, while having a large difference in the other channel. For instance, the top and bottom vectors differ slightly in the channel 1 value whereas they exhibit a large difference in the channel 2. When considering each channel separately, these vectors tend to have only their channel 1 averaged, while no averaging will be made when considering both channels simultaneously. As a result, the RMSE for the joint denoising reached lower values.

3.3. Denoising a multi-channel data with a non-linear target transformation

Some signals of interest are non-linear transforms and/or reductions from a higher-dimensional vectors to a lower-dimensional ones. In these cases, a variance-stabilizing transform might not be possible since the variance might not be a function of the mean.

As an example, consider a simulation of fluorescence anisotropy measurements. Light intensity was measured in two channels, parallel

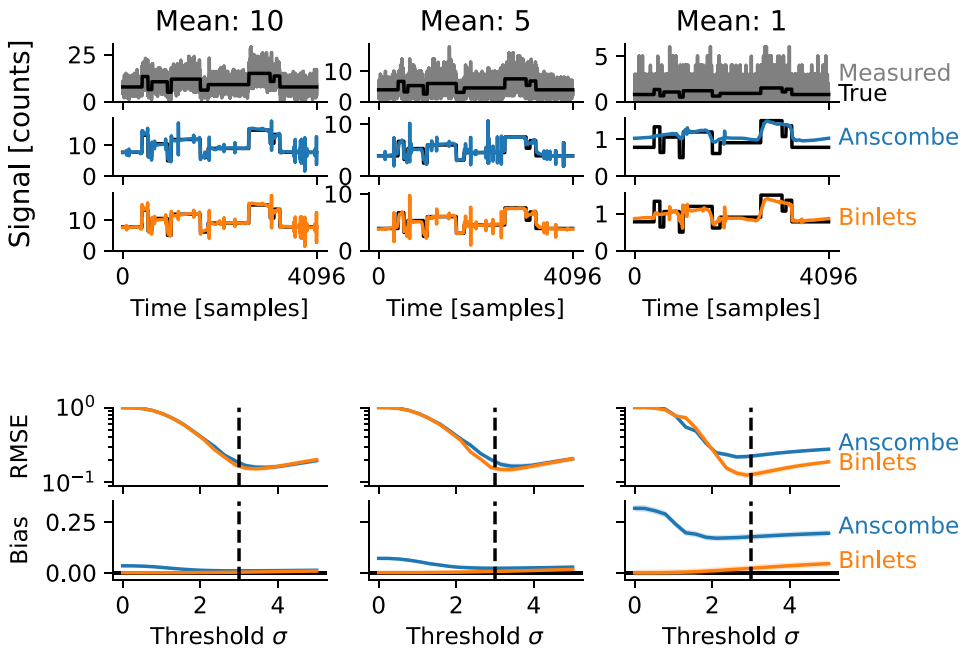

Fig. 2. Sampling from a Poisson distribution with a mean given by the true signal (black) produces a measured signal (gray), which is denoised via standard wavelet thresholding after applying an Anscombe transform (blue) and binlets with a Z-score test (orange). The chosen threshold σ is indicated by a black vertical line. For both methods, the RMSE and bias were calculated as functions of the threshold for several measurements. Binlets achieves a 1.75x improvement in the RMSE value in the last case, relative to the Anscombe transform-based denoising.

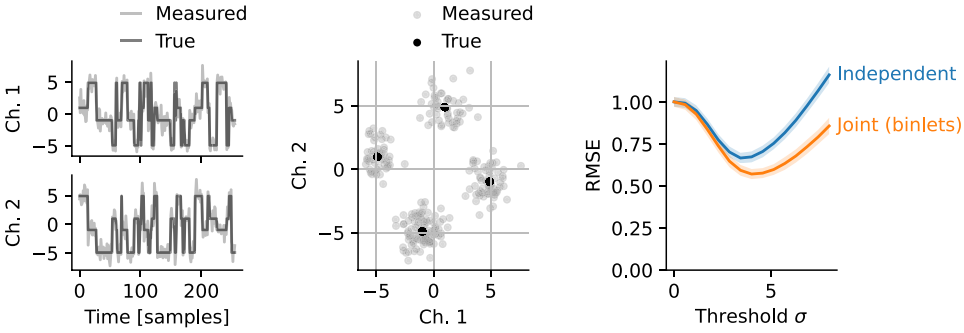

Fig. 3. A two-channel true signal (black line) was generated from four true vectors (black dots) by randomly alternating between them after an exponentially-distributed delay. The measured signal (gray) was generated by adding standard normal noise. The RMSE is calculated jointly for both channels as a function of the denoising threshold for two methods: independent, which analyzes each channel separately, and joint, which analyzes both channels jointly.

$I_\parallel$ and perpendicular $I_\perp$. The total intensity I_T and anisotropy r was calculated as [22]:

$$\begin{cases} I_T &= I_\parallel + 2I_\perp \\ r &= (I_\parallel - I_\perp)/I_T. \end{cases} \tag{7}$$

Hence, we have a two-dimensional vector $[I_\parallel, I_\perp]$ which was reduced to a scalar r with a non-linear transform. A variance-stabilizing transform of the anisotropy signal r is not possible, as the same value of r can come from different intensity values and have different variances.

In Fig. 4, we start from the true anisotropy r and total intensity I_T signals, generate $[I_\parallel, I_\perp]$ by inverting (7), and add Poisson noise to each channel. We consider two cases in which the true anisotropy signal is given by the *blocks* signal [20]. In one case, the true total intensity was constant, generating two clean signals for the true parallel and perpendicular intensities. In the other case, the true total intensity was chosen by sampling from a uniform distribution, giving true parallel and perpendicular intensities which already seem to be noisy. Nevertheless, if we computed the anisotropy from these, that is, before adding Poisson noise, we would obtain the true anisotropy signal (black).

On the right, we compare three different approaches to denoising: the first corresponds to analyzing each channel separately, with a Z-score test for Poisson noise (4), as in Section 3.1. The second method thresholds both channels simultaneously, as in Section 3.2, considering a diagonal covariance matrix for independent Poisson noise in each channel. Lastly, we also consider both channels simultaneously, but perform the test to threshold the detail coefficients in the anisotropy space using binlets. That is, for each $[I_\parallel, I_\perp]$ vector, we compute the anisotropy r and its variance, using linear error propagation, and do a Z-score test.

In the case of constant total intensity, we recovered the result of the previous section, where analyzing both channels simultaneously improved the denoising. However, when the total intensity fluctuated, the third approach, comparing in the anisotropy space, outperformed the other two methods by a large margin. This was expected as two completely different intensity vectors, which might not be averaged based on their intensity values, could correspond to the same anisotropy value; hence, it would be correct to average them.

We extended this example to a two dimensional signal with the same anisotropy profile, but regions with different distributions for the total intensity (Fig. 5). Where the true total intensity is homogeneous,

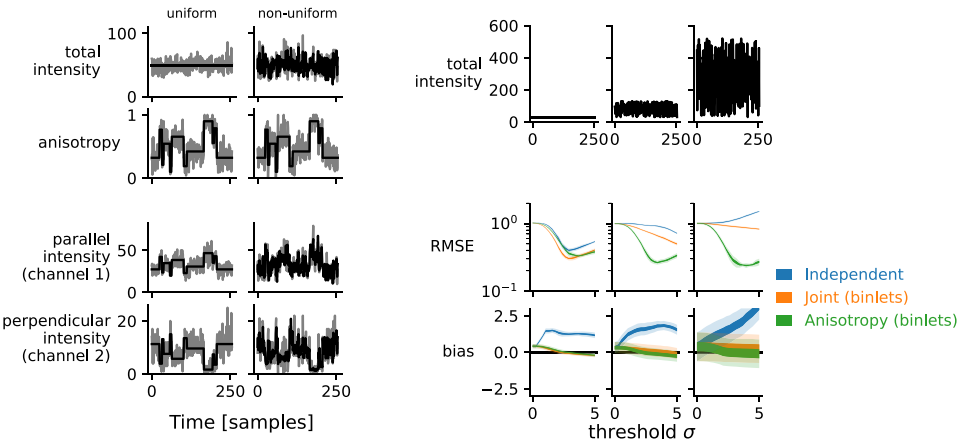

Fig. 4. Fluorescence anisotropy signals with different total intensity characteristics. (left) Two signals with the same true anisotropy profile (black), but homogeneous (first column) and non-homogeneous (second column) true total intensity (black). From these, the true parallel and perpendicular polarization intensity channels were computed, and the measurement (gray) simulated by adding Poisson noise. (right) RMSE and bias in the denoised anisotropy signal for three signals which have the same anisotropy pattern but different (true) total intensities.

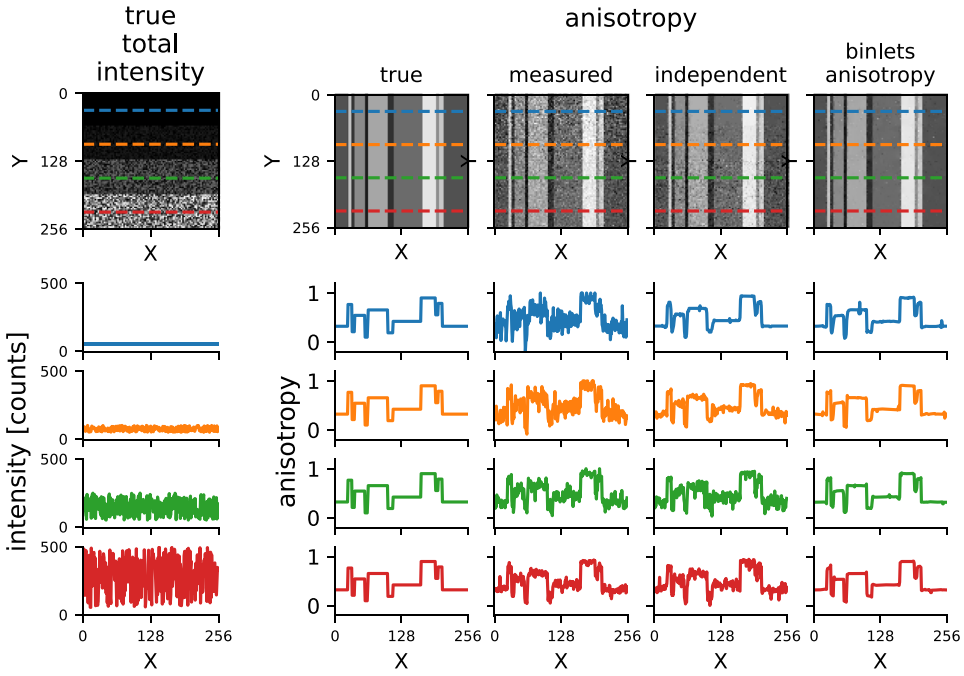

Fig. 5. Denoising of fluorescence anisotropy signals in two spatial dimensions. The top-left image corresponds to a simulated two-dimensional image of total intensity with a horizontal stripes pattern. It was divided into four regions, sampled from uniform distributions with different widths. Colored lines correspond to line profiles shown below. On its right, there are anisotropy images with a vertical stripes pattern, which correspond to the *blocks* [20] pattern used in the 1D case. From left to right, the columns correspond to the ground-truth, simulated measurement (no-denoising), independent denoising of each channel, and joint denoising with binlets.

both methods denoised the image correctly. However, in the noisier regions, only binlets was able to recognize the underlying regularity to correctly average neighboring pixels.

3.4. Denoising an experimental microscopy dataset

To test our method on experimental data, we used a publicly available dataset [17] consisting of fluorescence microscopy images obtained using different techniques. Each field of view was captured 50 times under the same instrumental settings, and a ground truth was constructed as the average of the images. In particular, we used the ones corresponding to confocal images of fixed bovine pulmonary artery endothelial (BPAE) cells labeled with DAPI to show the nucleus.

In [17], the authors trained two deep learning (DL) models, for which they made the weights available and we used to compare our method. One of the models is a feed-forward denoising convolutional neural network (DnCNN) [23], which requires pairs of noisy and ground truth images for testing. The other, Noise2Noise [18], does not require a ground truth and is trained from the noisy image itself.

While our method does not require training, it requires knowledge of the noise distribution. For microscopy images, it can be modeled as a Poisson distribution for the photon sensing, and an offset and scale factors due to the detector settings [24]. We used the 50 samples for the first field-of-view to fit this scale and offset.

Visually, our method showed a high resemblance to the ground-truth image (Fig. 6, left) and outperformed both deep learning methods. Considering the normalized residuals, binlets showed randomly distributed residuals, while both DL methods showed a bias towards lower intensity values: DnCNN in the nucleus, and N2N in the areas surrounding the nucleus and background.

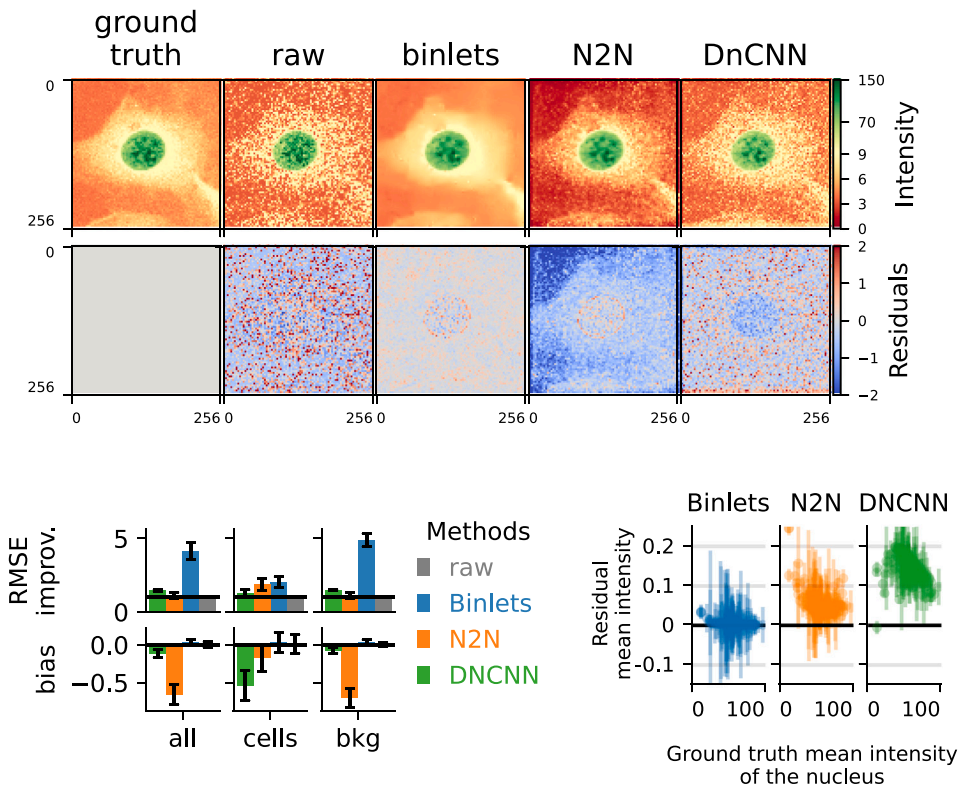

Fig. 6. Denoising applied to fluorescence microscopy images of cells. (left) Intensity images for the ground truth, raw, and raw denoised with binlets and two deep learning methods (N2N and DnCNN). The color scale changes slope at 9 to highlight differences in three areas: nucleus, the rest of the cell, and background. Residuals were calculated against the ground truth, which is the average of 50 raw images, and normalized with their standard deviation. (top-right) Mean and standard deviation for the metrics applied to the whole dataset: root mean squared error improvement (RMSE improv., normalized to raw) and mean bias. They were calculated in three regions: the whole image (all), background and cell areas. (bottom right) Residual mean intensity for each nucleus between the denoised and ground truth images as a function of the ground truth mean intensity. The error bars correspond to the standard deviation of the 50 samples.

For the entire dataset, we computed two summary metrics: the root mean squared error (RMSE) and bias (Fig. 6, top-right). These were computed for the whole image (all), but also separately for the background and cell regions. Our method shows equivalent or superior performance to the deep learning models, albeit with much less computational and experimental costs, as no training dataset or run is required. It must also be noted that most images were present in their training set and only one was from the testing set.

In practice, we are not interested in the overall error in the image, but we want to compute some observables of interest from it. A typical one in cell microscopy is the mean intensity in the nucleus. Binlets did not present any bias in the mean intensity for each nucleus (Fig. 6, bottom-right), in contrast to both deep learning methods.

4. Discussion

In this article, we presented a wavelet-based denoising method for single or multichannel signals of any dimensionality. It outperforms standard wavelet methods when the magnitude of interest involves a non-linear transformation or reduction of a multichannel signal. Examples of these signals include any normalized value, such as ratiometric observables, or more complex ones, such as phasor values in fluorescence lifetime imaging microscopy (FLIM) [14]. In these cases, the standard approach with wavelets can result in no denoising at all.

This method proposes to perform the denoising before the transformation is applied, but considering the intended transformation to discard coefficients. It presents a new thresholding method for detail coefficients, which takes the approximation coefficients used to compute them, applies the transformation, and compares them in that transformed space. Hence, averaging is performed in the original space, whereas the decision to average was made in the transformed space. Additionally, this method allows to leverage information in the correlation between channels. This might not be possible with the standard approach as a covariance-stabilizing transform might not exist [21].

Our method does not assume any temporal or spatial dimensionality in the signal. However, in the current implementation, it can only average square-shaped regions. As wavelet analysis has been extended to use basis functions that are localized in orientation, such as curvelets [10], future binlets extensions could overcome this limitation. Our method's only hypothesis is that neighboring data points might correspond to the same value, which is assessed by providing a statistical test to compare them, and can be averaged together. While it does not assume any type of noise distribution, it requires knowledge of the instrument-dependent noise distribution. However, once a calibration is performed, it can be used for any type of sample or signal. This is in stark contrast to *ad-hoc* deep learning *models*, which should not be reused for different samples [16].

Additionally, our method did not present any bias, either globally or in the regions of interest in the experimental dataset used. This was not the case for the deep learning methods considered, which might be biased towards the structures seen in the training process. Moreover, whereas deep learning methods are black boxes, our method offers a simple interpretation: a spatially-adaptive binning of the data. It provides a dial to adjust the trade-off between resolution and accuracy: the significance of the statistical test. Hence, it is easy to understand what it is doing and its limitations.

Finally, our method is simple to implement and fast without requiring any particular hardware such as GPUs. We provide a Python reference implementation which is available on GitHub (https://github.com/maurosilber/binlets) and installable from PyPI. It is easy to integrate into any pipeline analysis as it does not require any particular data format.

CRediT authorship contribution statement

Mauro Silberberg: Conceptualization, Software, Formal analysis, Writing – original draft, Investigation, Methodology. **Hernán E. Grecco:** Conceptualization, Investigation, Methodology, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data is available at https://github.com/maurosilber/binlets.

Acknowledgments

We thank Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET) and University of Buenos Aires (UBA) for financial support to M.S. and H.E.G. This work was supported by the following Grants: PICT 2014-3658.

Appendix. Pseudocode

Algorithm 1 Pseudocode for 1D binlets

function BINLETS(x, L, f)
Require: $x \in \mathbb{R}^{C\times N}$, $L \in \mathbb{N}$, $f : \mathbb{R}^C \times \mathbb{R}^C \to \{0, 1\}$
Ensure: $0 \le L^{max} \le \log_2(N)$

▷ **Decomposition and thresholding**
- $a_i^{L=0} \leftarrow x_i$ ▷ Input is the first approx. coeff.
- $m_i^{L=0} \leftarrow True$ ▷ Mask coefficients.
- **for** $L = 0 \dots L^{max}$ **do** ▷ For each level
  - **for** $i = 0 \dots N$ **do** ▷ For each time-point
    - $s \leftarrow 2^L$ ▷ Temporal shift
    - $a_i^{L+1} \leftarrow a_{i+s}^L + a_i^L$ ▷ Approx. coeff.
    - $d_i^{L+1} \leftarrow a_{i+s}^L - a_i^L$ ▷ Detail coeff.
    - ▷ Up to here, standard Haar decomposition
    - ▷ Now, test based on comparison of approx. coefficients
    - **if** $f(a_i^L, a_{i+s}^L)$ **and** m_i **then** ▷ Compare approx from previous level L
      - $d_i^{L+1} \leftarrow 0$ ▷ Discard detail coefficient of level $L + 1$.
    - **else**
      - $m_i \leftarrow False$
    - **end if**
  - **end for**
- **end for**

▷ **Reconstruction**
- **for** $L = L^{max} \dots 0$ **do** ▷ For each level in reverse order
  - **for** $i = 0 \dots N$ **do** ▷ For each time-point
    - $s \leftarrow 2^L$
    - $a_i^{L-1} \leftarrow (a_{i+s}^L - d_i^L)/2$
    - $a_{i+s}^{L-1} \leftarrow (a_{i+s}^L + d_i^L)/2$
  - **end for**
- **end for**

- **return** $a_{L=0}$

end function

References

[1] F.J. Anscombe, Graphs in statistical analysis, Amer. Statist. 27 (1) (1973) 17–21, http://dx.doi.org/10.1080/00031305.1973.10478966.
[2] W. Meiniel, J.-C. Olivo-Marin, E.D. Angelini, Denoising of microscopy images: A review of the state-of-the-art, and a new sparsity-based method, IEEE Trans. Image Process. 27 (8) (2018) 3842–3856, http://dx.doi.org/10.1109/TIP.2018.2819821.
[3] L. Fan, F. Zhang, H. Fan, C. Zhang, Brief review of image denoising techniques, Vis. Comput. Ind. Biomed. Art 2 (1) (2019) 7, http://dx.doi.org/10.1186/s42492-019-0016-7.
[4] B. Goyal, A. Dogra, S. Agrawal, B.S. Sohi, A. Sharma, Image denoising review: From classical to state-of-the-art approaches, Inf. Fusion 55 (2020) 220–244, http://dx.doi.org/10.1016/j.inffus.2019.09.003.
[5] G. Chen, W. Xie, Y. Zhao, Wavelet-based denoising: A brief review, in: 2013 Fourth International Conference on Intelligent Control and Information Processing, ICICIP, 2013, pp. 570–574, http://dx.doi.org/10.1109/ICICIP.2013.6568140.
[6] A. Cohen, M.A. Atoui, On wavelet-based statistical process monitoring, Trans. Inst. Meas. Control 44 (3) (2022) 525–538, http://dx.doi.org/10.1177/0142331220935708.
[7] M. Makitalo, A. Foi, A closed-form approximation of the exact unbiased inverse of the anscombe variance-stabilizing transformation, IEEE Trans. Image Process. 20 (9) (2011) 2697–2698, http://dx.doi.org/10.1109/TIP.2011.2121085.
[8] B. Zhang, M.J. Fadili, J.L. Starck, S.W. Digel, Fast poisson noise removal by biorthogonal haar domain hypothesis testing, Stat. Methodol. 5 (4) (2008) 387–396, http://dx.doi.org/10.1016/j.stamet.2008.02.004.
[9] K. Hirakawa, P.J. Wolfe, Skellam shrinkage: Wavelet-based intensity estimation for inhomogeneous poisson data, IEEE Trans. Inform. Theory 58 (2) (2012) 1080–1093, http://dx.doi.org/10.1109/TIT.2011.2165933.
[10] S.G. Mallat, A Wavelet Tour of Signal Processing: The Sparse Way, third ed., Elsevier/Academic Press, Amsterdam ; Boston, 2009.
[11] X.-g. Xia, B. Suter, Vector-valued wavelets and vector filter banks, IEEE Trans. Signal Process. 44 (3) (1996) 508–518, http://dx.doi.org/10.1109/78.489024.
[12] M.A. Westenberg, T. Ertl, Denoising 2-D Vector Fields by Vector Wavelet Thresholding, Václav Skala - UNION Agency, 2005.
[13] B. Spring, R. Clegg, Image analysis for denoising full-field frequency-domain fluorescence lifetime images, J. Microsc. 235 (2) (2009) 221–237, http://dx.doi.org/10.1111/j.1365-2818.2009.03212.x.
[14] M. Silberberg, H.E. Grecco, pawFLIM: Reducing bias and uncertainty to enable lower photon count in FLIM experiments, Methods Appl. Fluoresc. 5 (2) (2017) 024016, http://dx.doi.org/10.1088/2050-6120/aa72ab.
[15] R.F. Laine, G. Jacquemet, A. Krull, Imaging in focus: An introduction to denoising bioimages in the era of deep learning, Int. J. Biochem. Cell Biol. 140 (2021) 106077, http://dx.doi.org/10.1016/j.biocel.2021.106077.
[16] CSBDeep - a deep learning toolbox for micrsocopy image restoration and analysis - Frequently asked questions, 2023, CSBDeep https://csbdeep.bioimagecomputing.com/faq.
[17] Y. Zhang, Y. Zhu, E. Nichols, Q. Wang, S. Zhang, C. Smith, S. Howard, A Poisson-Gaussian denoising dataset with real fluorescence microscopy images, in: CVPR, 2019, http://dx.doi.org/10.48550/arXiv.1812.10366.
[18] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, T. Aila, Noise2Noise: Learning image restoration without clean data, 2018, http://dx.doi.org/10.48550/ARXIV.1803.04189, arXiv.
[19] D.L. Cade, Gigapixel AI accidentally added ryan gosling's face to this photo, 2020, PetaPixel https://petapixel.com/2020/08/17/gigapixel-ai-accidentally-added-ryan-goslings-face-to-this-photo/.
[20] D.L. Donoho, I.M. Johnstone, Ideal spatial adaptation by wavelet shrinkage, Biometrika 81 (3) (1994) 425–455, http://dx.doi.org/10.1093/biomet/81.3.425.
[21] P.W. Holland, Covariance stabilizing transformations, Ann. Statist. 1 (1) (1973) 84–92, arXiv:2958159.
[22] J.R. Lakowicz (Ed.), Principles of fluorescence spectroscopy, Springer US, Boston, MA, 2006, http://dx.doi.org/10.1007/978-0-387-46312-4.
[23] K. Zhang, W. Zuo, Y. Chen, D. Meng, L. Zhang, Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising, IEEE Trans. Image Process. 26 (7) (2017) 3142–3155, http://dx.doi.org/10.1109/TIP.2017.2662206.
[24] A. Foi, M. Trimeche, V. Katkovnik, K. Egiazarian, Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data, IEEE Trans. Image Process. 17 (10) (2008) 1737–1754, http://dx.doi.org/10.1109/TIP.2008.2001399.