



Filtering Quasiperiodic Noise from Biomedical Images

Máté Bakos

Master of Science Thesis

Filtering Quasiperiodic Noise from Biomedical Images

MASTER OF SCIENCE THESIS

For the degree of Master of Science in Systems and Control at Delft
University of Technology

Máté Bakos

Friday 17th July, 2020

DELFT UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF
DELFT CENTER FOR SYSTEMS AND CONTROL (DCSC)

The undersigned hereby certify that they have read and recommend to the Faculty
of Mechanical, Maritime and Materials Engineering (3mE) for acceptance a thesis
entitled

FILTERING QUASIPERIODIC NOISE FROM BIOMEDICAL IMAGES

by

MÁTÉ BAKOS

in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE SYSTEMS AND CONTROL

Dated: Friday 17th July, 2020

Supervisor(s):

dr.eng. Raf Van de Plas

Reader(s):

dr.eng. Sander Wahls

dr. Matthias Alfeld

Abstract

In the field of biomedical imaging, images often report a combination of biologically induced variation, usually the goal of the imaging process (e.g. outlining an anatomical region or disease pattern), and non-biological variation, such as instrument or acquisition method-induced noise patterns. Since some medical decisions are made based on imaging, separating the biological signal from noise is of significant importance (e.g. accelerates decision-making, reducing the chance of misdiagnosis). Some non-biological variations that span a wide range of imaging modalities include e.g. viewport stitching artifacts, slice-to-slice interference, aliasing, and Gibbs-phenomena. From a signal processing perspective, many of these can be modeled as quasiperiodic patterns. Thus, removal of quasiperiodic patterns while preserving the underlying medical information is the main focus of this thesis.

Although in modern instruments, many forms of non-biological variation can be attenuated to be invisible to the naked eye, machine learning algorithms which are often used for classification of disease and segmentation of biological samples may be susceptible to even minor variations and noise patterns. Development of entirely data-driven, unsupervised denoising techniques can potentially increase the effectiveness and reliability of such algorithms. Furthermore, under certain image transformations, such as different color spaces, the Fourier and wavelet transform, and factorizations, such as principal component analysis and non-negative matrix factorization, as well as combinations of these, non-biological patterns can get amplified and become so prominent that much of the underlying biological information is concealed. Removing these quasiperiodic patterns using current state-of-the-art algorithms still requires manual parameter tuning and prior expert knowledge, which is an impractical and possibly unnecessary expectation towards healthcare professionals. The goal of this M.Sc. thesis is to develop an automated, data-driven framework, that is able to reliably identify, quantify, and eliminate quasiperiodic patterns within the images while retaining as much biological information as possible.

In this framework, named Quasiperiodic Image Denoising (QID), two novel algorithms are implemented, both operating in the Fourier domain. One algorithm is based on robust principal component analysis (QID-RPCA) and the other uses the normalized median of absolute differences (QID-MADN). The methods used to achieve unsupervised, data-driven denoising are described in detail. This includes the use of histogram equalization for radial binning, an automated, sparsity-based approach to choosing the optimal aggregation level, and noise component attenuation based on radial frequency patterns. The methodology is demonstrated through three case studies. First, a synthetic dataset is used to compare the performance of the novel algorithms to the current state-of-the-art solutions. Second, the performance is evaluated on two real-world datasets processed using a number of methods, e.g. factorization and different color spaces. One of these datasets is

based on a microscopy image of a transversal section of a mouse brain and the other one is based on a microscopy image of a coronal section of a rat kidney. Finally, a real-world, raw dataset is denoised consisting of a set of high-resolution fluorescent microscopy images of a human kidney. Results indicate that the novel algorithms have higher denoising performance than previous approaches in the literature with notable improvements achieved for low-frequency corruptions.

Contents

Abstract	i
List of Figures	ix
List of Tables	xiii
Acknowledgements	xv
1 Introduction	1
2 Background and Fundamentals	5
2-1 Medical imaging	5
2-1-1 Microscopy	6
2-1-2 Other imaging modalities	8
2-2 Common noise types and removal	9
2-2-1 Gaussian noise	10
2-2-2 Impulse or salt-and-pepper noise	11
2-2-3 Periodic and quasiperiodic noise	12
2-2-4 Structured noise	14
2-3 Decomposition and domain transformation	15
2-3-1 Signal domains	15
2-3-1-1 Spatial-domain parametric models	15
2-3-1-2 Fourier-transform and frequency-domain analysis	16
2-3-1-3 Wavelet-transform	19
2-3-2 Matrix factorization methods	20
2-3-2-1 Principal component analysis	20
2-3-2-2 Non-negative matrix factorization	22
2-3-2-3 Robust principal component analysis	22
2-4 Filtering and identification	22
2-4-1 Recognition and clustering of outliers	23
2-4-1-1 Z-score and three-sigma rule	23
2-4-1-2 Modified Z-score or robust three-sigma rule	24
2-4-1-3 Interquartile range	24
2-4-1-4 Hard and soft thresholding	25
2-4-2 Filtering	26

2-4-2-1	Adaptive thresholding	26
2-4-2-2	Frequency domain filters	27
2-4-2-3	Noise component attenuation	28
2-5	Image restoration	29
2-5-1	Total variation minimization	29
2-5-2	Metrics	30
2-5-2-1	Signal-to-noise ratio	30
2-5-2-2	Peak signal-to-noise ratio	30
2-5-2-3	Structural similarity index	31
2-5-2-4	Blind/Referenceless Image Spatial Quality Evaluator	31
2-6	State-of-the-art	31
2-6-1	ARPENOS	32
2-6-2	ACARPENOS	34
2-6-3	IONITA	35
2-7	Summary	37
3	Methods	41
3-1	Pipeline	43
3-2	Apodization and padding	45
3-3	Fourier transform	48
3-4	Protection area for the zero frequency and adjacent components	50
3-5	Cartesian to polar	52
3-6	Radial binning with histogram equalization	53
3-7	Mapping polar values to a rectangular grid	57
3-8	Robust Principal Component Analysis	60
3-9	Outlier detection	61
3-10	Optimization of the binning strategy	63
3-11	Noise component attenuation	66
3-12	Image reconstruction by inverse Fourier transform	67
4	Case studies	69
4-1	Shepp-Logan synthetic dataset	69
4-1-1	Experimental setup	70
4-2	Processed mouse brain and rat kidney dataset	72
4-2-1	Experimental setup	72
4-3	Raw human kidney dataset	73
4-3-1	Experimental setup	74
5	Results and Discussion	77
5-1	Synthetic dataset results and performance comparison	77
5-1-1	Frequency domain performance comparison of QID-RPCA and QID-MADN versus prior algorithms from the literature	78
5-1-2	Image domain performance comparison of QID versus the state-of-the-art algorithms from literature	84
5-1-3	Information loss through saturation of the dynamic range	89
5-1-4	Performance for denoising corruptions over the Nyquist-frequency	92
5-1-5	Runtime scaling	95
5-2	Processed dataset results and performance comparison	95
5-3	Real-world microscopy dataset results and performance comparison	100
5-3-1	EGFP band, in-depth analysis	105
5-3-2	DAPI band, in-depth analysis	110
5-3-3	Cy5 band, in-depth analysis	112

5-3-4 Summary	116
6 Conclusions and future work	119
6-1 Conclusions	119
6-2 Future work	120
Bibliography	123

List of Abbreviations

Abbreviations

2-D	Two-dimensional
ADC	Analogue-to-digital converter
BRISQUE	Blind/Referenceless Image Spatial Quality Evaluator
CFT	Continuous Fourier Transform
CT	Computer Tomography
dB	Decibel
DC	Direct Current
DFT	Discrete Fourier Transform
FFT	Fast Fourier Transform
IMS	Imaging Mass Spectrometry
IQR	Interquartile Range
LUD	Lower-upper Decomposition
MADN	Median Absolute Deviation Normalized
MRI	Magnetic Resonance Imaging
MSE	Mean Squared Error
NFA	Number of False Alarms
NNMF	Non-negative Matrix Factorization
NSS	Natural scene statistics
PC	Principal Component
PCA	Principal Component Analysis
PDF	Probability Density Function
PET	Positron Emission Tomography
PSNR	Peak Signal-to-Noise Ratio
QID	Quasiperiodic Image Denoising
RMSE	Root Mean Squared Error
RPCA	Robust Principal Component Analysis
SEM	Scanning Electron Microscopy
SNR	Signal-to-Noise Ratio
SOTA	State-of-the-art
SSIM	Structural Similarity index
SURE	Stein's Unbiased Risk Estimate
TEM	Transmission Electron Microscopy

List of Figures

1-1	Example of a biomedical image with and without acquisition method-induced quasiperiodic noise.	3
2-1	Overview of the electromagnetic spectrum showing energy per photon, frequency, wavelength, and their scale as well as the radiation types. Source: Reproduced from [85], under Creative Commons Attribution-ShareAlike License.	6
2-2	Examples of specialized techniques in microscopy. Source: [68], Figure 3, under Creative Commons Attribution License 3.0. Source: [91], under Creative Commons Attribution License 4.0. Source: [87], under Creative Commons License 1.0.	7
2-3	Examples of various noise artefacts on a variety of biomedical images [35], [59]. Source: [35], Figures 7,8,9,11, under Creative Commons Attribution License 4.0. Source: [59], Figure 17,25.	9
2-4	An example of the effect of Gaussian noise.	11
2-5	An example of the effect of impulse noise.	12
2-6	An example of the effect of periodic noise.	13
2-7	Examples of Moiré patterns. Source: [86], Figure 1.1, reproduced with permission from Springer-Verlag London Limited.	14
2-8	Viewport-stitching artifact on H&E stained microscopy images. Source: [56], Figure 5, under Creative Commons Attribution License 4.0.	15
2-9	The sampled version of two different signals is identical. The higher-frequency signal shown on Figure 2-9 is sampled insufficiently, as the discrete waveform is not unique, while the lower-frequency signal is. Source: [64], Figure 5.1, reproduced with permission from Springer-Verlag London Limited.	17
2-10	Examples of the effect of finite windowing both in the signal and frequency domain. Source: [64], Figure 6.2, reproduced with permission from Springer-Verlag London Limited.	19
2-11	Examples of leakage both in the signal and frequency domain. Source: [64], Figure 6.3, reproduced with permission from Springer-Verlag London Limited.	19
2-12	Two-dimensional example of coordinate system transformation using PCA. PC1 captures the majority of variance in the dataset. While PC2 still captures some variance, it is considerably less, compared to PC1 [90].	21
2-13	Wavelet thresholding methods.	26
2-14	The effect of different basic filtering methods in the Fourier-domain.	28

2-15	Main steps of the ARPENOS algorithm [74]. In the one-dimensional image, the red line indicates the regression line, its horizontal limits indicate the borders of the linear region. The green line indicates the upper bound that is three standard deviations over the trendline, horizontally, it shows the lower-frequency detection limit to be at 0.04 cycles/pixel. Source: [74], Figures 16,17, rights to reuse requested from the Society of PhotoOptical Instrumentation Engineers (SPIE).	33
2-16	Main steps of the ACARPENOS algorithm. In the outlier map yellow indicates outliers, which are an unexpectedly high NFAs. The central area is protected from filtering. Source: [72], Figure 1, ©2015 IEEE.	35
2-17	Main steps of the IONITA algorithm [70]. The specific threshold used, the outlier map, and the scales of the power spectrum are not shown as they were not provided in the original paper. Source: [70], Figure 2,4,6, ©2015 IEEE. . .	37
3-1	Examples of the most important steps of the QID denoising pipeline.	42
3-2	Flowchart detailing the steps of each version of the algorithm.	43
3-3	Pipeline blocks relevant to padding and apodization.	45
3-4	An example of the Gibbs-phenomenon. It appears at discontinuities of the time or image domain when the discrete Fourier transform approximates the theoretical Fourier representation with a limited range.	45
3-5	Examples of the effects of padding and windowing. Highlighting the emergence of artificial patterns.	48
3-6	Examples of the Fourier transformed values of a corrupt image.	49
3-7	Pipeline blocks relevant to protection area masking.	50
3-8	Examples of the effect of the protection area radius.	51
3-9	Pipeline blocks relevant to polar coordinate system transformation.	52
3-10	Examples of the visual representation of the polar coordinates of a power spectrum.	52
3-11	Pipeline blocks relevant to radial histogram equalization and binning.	53
3-12	An example of increasing the contrast of an image through histogram equalization of intensity values.	54
3-13	Examples of equidistant and histogram equalized radial binning from the central coefficient. With equidistant binning, bins close to the center contain a low number of coefficients, while bins farther from the center contain an increasingly higher number of coefficients. The histogram equalized binning strategy, contains an approximately equal number of coefficient values in each bin. Both figures show a zoomed-in, central section for visual clarity. The color coding indicates the index of the radial bin that the coefficient gets sorted into.	55
3-14	Examples of common-radius masking used for QID-RPCA and omitted for QID-MADN. Both histogram equalized binning strategies consist of 24 distinct bins.	56
3-15	Pipeline blocks relevant to mapping polar values to a rectangular grid.	57
3-16	Example of combining radial and angular binning for the rectangular representation of polar values.	59
3-17	Example of the logarithm of the power spectrum, aggregated to the rectangular grid of polar values.	59
3-18	Pipeline blocks relevant to robust principal component analysis.	60
3-19	Examples of the input and outputs of RPCA. The logarithm of the power spectrum in a rectangular, polar coordinate matrix M is separated into its low-rank component, L , and sparse component, S	61
3-20	Pipeline blocks relevant to outlier labeling.	61

3-21	Examples of the outlier maps of QID-RPCA and QID-MADN. The outlier map of QID-RPCA is shown both in the rectangular, polar coordinate-system, as well as the original Cartesian coordinate-system.	63
3-22	Pipeline blocks relevant to noise component attenuation.	66
3-23	Comparison of the original and attenuated power spectrums.	66
3-24	Pipeline blocks relevant to transform inversion.	67
3-25	Comparison of the original and denoised images.	68
4-1	Modified Shepp-Logan phantom.	70
4-2	Processed microscopy images, on a false color scale. Processing includes different color spaces such as RGB, HVS, and YCbCr, specific factors of techniques such as PCA, NMF, and DWT, range- and median filtered, entropy- and median filtered versions.	72
4-3	EGFP band of the human kidney microscopy dataset. Note the grid-like repetitive pattern resulting from viewport being stitched together, superimposed on top of the biological pattern of the human kidney.	74
5-1	Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of MSE score.	79
5-2	Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of MSE score.	79
5-3	Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of SSIM score.	80
5-4	Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of SSIM score.	80
5-5	Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of SNR score.	81
5-6	Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of SNR score.	81
5-7	Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of PSNR score.	82
5-8	Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of PSNR score.	82
5-9	Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of BRISQUE score.	83
5-10	Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of BRISQUE score.	83
5-11	Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.06 cycles/pixel, noise amplitude = 25 %. . .	84
5-12	Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.08 cycles/pixel, noise amplitude = 25 %. . .	85
5-13	Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.02 cycles/pixel, noise amplitude = 25 %. . .	86
5-14	Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.14 cycles/pixel, noise amplitude = 25 %. . .	87
5-15	Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.5 cycles/pixel, noise amplitude = 25 %. . .	88
5-16	Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.04 cycles/pixel, noise amplitude = 5 %. . .	89
5-17	Effect of information loss through the saturation of the dynamic range of the image measured in terms of SSIM scores.	90

5-18	Effect of information loss through the saturation of the dynamic range of the image measured in terms of SSIM scores.	91
5-19	Comparison of the current State-of-the-art algorithms and QID under information loss. Noise frequency = 0.1 cycles/pixel, noise amplitude = 25 %.	91
5-20	Effect of noise frequencies over and under cutoff frequency.	93
5-21	Effect of noise frequencies over and under cutoff frequency.	93
5-22	Comparison of the current State-of-the-art algorithms and QID for corruption over the Nyquist-frequency. Noise frequency = 0.6 cycles/pixel, noise amplitude = 25 %.	94
5-23	Plots showing computation times of each algorithm as a function of number of pixels.	95
5-24	Noise removal of band 99 of the mouse brain dataset.	98
5-25	Noise removal of band 279 of the mouse brain dataset.	99
5-26	Noise removal of band 99 of the rat kidney dataset.	99
5-27	Noise removal of band 309 of the rat kidney dataset.	100
5-28	EGFP band of the human kidney fluorescent microscopy dataset.	101
5-29	DAPI band of the human kidney fluorescent microscopy dataset.	102
5-30	Cy5 band of the human kidney fluorescent microscopy dataset.	103
5-31	DsRed band of the human kidney fluorescent microscopy dataset.	104
5-32	Noise removal of EGFP band of the human kidney fluorescent microscopy dataset.	105
5-33	Noise removal of EGFP band of the human kidney fluorescent microscopy dataset.	107
5-34	Zoomed-in portion of the central part of the EGFP band of the human kidney fluorescent microscopy dataset.	108
5-35	Zoomed-in portion of the right corner of the EGFP band of the human kidney fluorescent microscopy dataset.	109
5-36	Zoomed-in central portion of the EGFP band of the logarithm of the raw power spectrum, the outlier map, and the logarithm of the filtered power spectrum.	110
5-37	Noise removal of DAPI band of the human kidney fluorescent microscopy dataset.	111
5-38	Zoomed-in central portion of the DAPI band of the logarithm of the raw power spectrum, the outlier map, and the logarithm of the filtered power spectrum.	112
5-39	Noise removal of Cy5 band of the human kidney fluorescent microscopy dataset.	113
5-40	Zoomed in portion of the central part of the Cy5 band of the human kidney fluorescent microscopy dataset.	114
5-41	Zoomed in portion of the right corner of the Cy5 band of the human kidney fluorescent microscopy dataset.	115
5-42	Zoomed-in central portion of the Cy5 band of the logarithm of the raw power spectrum, the outlier map, and the logarithm of the filtered power spectrum.	116

List of Tables

4-1

Table of filter and immunomarker information for the human kidney immunofluorescence image. Each row describes one of the four band examined in this experiment.

73

5-1

BRISQUE scores of selected bands from the mouse brain dataset and images denoised by ARPENOS, ACARPENOS, IONITA, QID-RPCA, and QID-MADN.

97

5-2

BRISQUE scores of selected bands from the rat kidney dataset and images denoised by ARPENOS, ACARPENOS, IONITA, QID-RPCA, and QID-MADN.

97

Acknowledgements

This thesis, as well as my master studies have been a wonderful experience. I believe, that I have matured a considerable amount, grew professionally and personally. I'd like to thank those, who supported me and without whom, this graduation would probably not have been possible.

First, I would like to thank my supervisor, Raf Van de Plas, for helping me through this thesis project. Thank you for your guidance, the intriguing conversations, and all I've learnt. You supported me and kept my hopes up constantly.

Next, I wish to thank all my friends in the Netherlands and my colleagues at Dutch Analytics. Special thanks to Kinga, my housemates, Simona, Kyle, and Laura, my classmates, Julia, Dan  l, Daniel, Ola, Karol, Leo, Patrick, Bart, Thomas, Jeroen, and Marc. You made my stay in Delft fun through rain and shine, I knew I could always count on you, you became a family to me that I can always come home to, wherever we are in the world. You taught me a lot about different cultures, habits, and perspectives on life. You taught me invaluable lessons, in your own way and helped me develop as a person. Thank you.

Then, I want to thank my family for always supporting me through good and bad times and being by my side.

Finally, I would like to thank Nathan Heath Patterson, Ph.D. of Vanderbilt University for providing the fluorescent microscopy dataset used for this thesis, and Pawe   Czerwi  ski for the cover page of this thesis.

It has been a wonderful experience overall, I have fond memories of this period and will certainly think back to it with a smile. I am grateful for all of you. Thank you!

Delft, University of Technology
Friday 17th July, 2020

M  t   Bakos

Chapter 1

Introduction

Medical and biological imaging are areas that acquire spatial measurements from organs, tissues, fluids, and other biological specimens. These imaging techniques are tools that can be used to achieve a wide range of goals. The most common goal of biomedical imaging is to conduct clinical analysis on the basis of images, including detection of disease, localization of abnormalities, and tracking their development [66]. There are also applications focusing on research and development, where images are used to gather knowledge on molecular pathomechanisms. Overall, imaging techniques serve as an essential step for many of today's prognostics and diagnostics.

Biomedical images often report a combination of biologically induced variation and non-biological variation, such as instrument or acquisition method-induced noise patterns [81], [94]. Some non-biological variations that span a wide range of imaging modalities include e.g. viewport stitching artifacts, slice-to-slice interference, aliasing, Gibbs-phenomenon, changing intensity gradients, ripple-effect, and wraparound of the sample [27], [35], [47], [59]. Some biological variations, such as membranes surrounding biological structures, cell boundaries, bones, and area specific tissue patterns may resemble non-biological variation [66]. Characterization and separation of biological and non-biological variations is often far from trivial [39], [45]. Many professionals in the life sciences, such as surgeons, radiologists, pathologists, bioinformaticians, and biomedical engineers handle such data day-to-day. Even though the use is wide-spread, it is common to see imaging measurements that are evidently corrupted by non-biological signals [35], [59]. This is not only bothersome, but it is also not part of the sample at hand, potentially obscuring important biological information and making discovery difficult or in some cases impossible. These phenomena can also be observed in research, where images containing a significant amount of noise are sometimes published. Since some medical decisions are made based on imaging, separating the biological signal from the non-biological variations is of significant importance. The most important difference to most fields that require denoising of images is that in biomedical imaging, introduction of false, artificial patterns can potentially have detrimental effects. Such patterns can be the difference between vastly different diagnoses. Introduced patterns have the potential to obstruct important biological structures, thus protection of the true biological signal is of paramount importance [45]. Sometimes preservation of the biological texture is more vital than removal of the noise.


Some of the noise sources causing the corruptions are well-understood mathematically, which makes characterization and removal possible [64], [86]. From the signal processing

perspective, many of these can be modeled as quasiperiodic patterns [69], [74], [77]. However, most of the popular approaches to such noise include tuning parameters that have to be changed manually and with every use, making wide-range adoption of such methods impractical and cumbersome without expert knowledge of the specific algorithm [32], [70], [72], [74], [82], [84]. By taking care of non-biological corruptions posterior to the imaging in an automatic, data-driven way, the accuracy of the results could be increased, previously invisible information could be revealed and it could enable a decrease in the chance of misdiagnosis. There is a wide variety of denoising algorithms focused on biomedical imaging [33], [36], [44], [46], [50], [62], [71], [76] and there is a wide variety of denoising algorithms focused on quasiperiodic patterns [30], [49], [70], [72], [74], [77], [84], however, the intersection of these fields is scarce.


It is important to note that in an ideal scenario these denoising approaches would be applied right after acquisition in the acquisition computer. Expert knowledge of digital signal processing would be encoded in the algorithms. The manual parameter setting would be avoided with data-driven algorithm design. Descriptive statistical properties of the biological signal and noise would be gathered automatically. No additional hardware should be required except for a general-purpose computer, making the approach both practical and economical.

This thesis explores what the currently applied, best-performing techniques for characterization and removal of periodic noise in images are, with particular focus on medical applications. The goal of this thesis is to develop a novel algorithmic framework that is able to reliably identify, quantify, and eliminate quasiperiodic, non-biological patterns within the images while retaining as much biological information as possible, in an automated, entirely data-driven fashion. An example of a biomedical image corrupted with acquisition method-induced quasiperiodic noise is shown in Figure 1-1a. An example of successful removal of such a quasiperiodic pattern, while preserving the underlying biological information is shown in Figure 1-1b.

The main questions this thesis focuses on are as follows. Applied on biomedical images, is it possible to surpass the performance of current state-of-the-art quasiperiodic denoising algorithms? Is it possible to create a reliable, data-driven framework for quasiperiodic denoising algorithms?



(a) Microscopy image with quasiperiodic corruption.



(b) Denoised microscopy image with no quasiperiodic corruption.

Figure 1-1: Example of a biomedical image with and without acquisition method-induced quasiperiodic noise.

Chapter 2 focuses on providing a background to the most common noise types in the medical domain. Furthermore, it aims to give an overview of the most commonly and successfully used techniques in a generic denoising pipeline with special focus on quasiperiodic patterns. Finally, it provides a brief introduction to some of the state-of-the-art algorithms in terms of removal of periodic and quasiperiodic corruptions and motivates the need for further development of these algorithms.

In Chapter 3, the novel framework operating in the Fourier domain is introduced, named Quasiperiodic Image Denoising (QID), and the steps of the two denoising algorithms developed as a part of this thesis, QID-RPCA and QID-MADN, are demonstrated and explained in detail. These steps build on previous developments and aim to overcome

their shortcomings, largely focusing on high denoising performance without the need for manual tuning of parameters. This includes the use of histogram equalization for radial binning; an automated, sparsity-based approach to choosing the optimal aggregation level; robustification of outlier detection through robust principal component analysis and robust median based statistical filtering; and finally, noise component attenuation based on radial frequency patterns.

Chapter 4 introduces three datasets and a set of experiments to demonstrate and validate the effectiveness of the algorithms. This includes a synthetic dataset and two, real-world microscopy datasets. Details of the datasets and noise artefacts are introduced and the experimental setup is outlined.

Chapter 5 applies the state-of-the-art and the novel algorithms to the case study images and discusses differences, strengths, and weaknesses of each. The methodology is demonstrated through three case studies. First, the synthetic dataset is evaluated on a wide range of simple periodic corruptions with a multitude of frequencies and amplitudes. Special cases such as information loss through intensity saturation and denoising performance close to the Nyquist-frequency are evaluated. Comparison of the denoising performance of the novel algorithms to the current state-of-the-art is provided through a number of metrics. Visual examples of the algorithms are provided, artifacts, strengths, and weaknesses are discussed in detail. Second, the performance is evaluated on two real-world datasets processed using a number of methods, e.g. factorization and different color spaces. One of these datasets is based on a microscopy image of a transversal section of a mouse brain and the other one is based on a microscopy image of a coronal section of a rat kidney. Finally, a real-world, raw dataset is denoised consisting of a set of high-resolution fluorescent microscopy images of a human kidney. Visual comparison of the denoising performance is given. Success of the initial goals of the research is evaluated.

Finally, in Chapter 6, concluding and summarizing remarks are given, and some ideas for further research and development are provided.

Chapter 2

Background and Fundamentals

2-1. Medical imaging

In this study, the use of the term 'medical imaging' is to be regarded in the broad sense of including all techniques that acquire spatial representations of the internal aspects of bodies, organs, and tissues, both *in vitro* and *in vivo*. The goal of these spatial representations is often detecting disease, finding abnormalities, and developing insights about molecular pathomechanisms, among others [28], [34], [48], [66]–[68].

In the medical imaging domain, many commonly used imaging techniques are based on the electromagnetic energy spectrum. These techniques can be sorted by energy per photon, as shown in Figure 2-1 [85]. At the lower end, there are radio waves used in magnetic resonance imaging (MRI). Then, there are infrared, visible, and ultraviolet rays used for microscopy, followed by X-ray radiation in skeletal system imaging and soft tissue disease. At the highest end of the spectrum, Gamma-rays are used for bone scans and positron emission tomography (PET) [81].

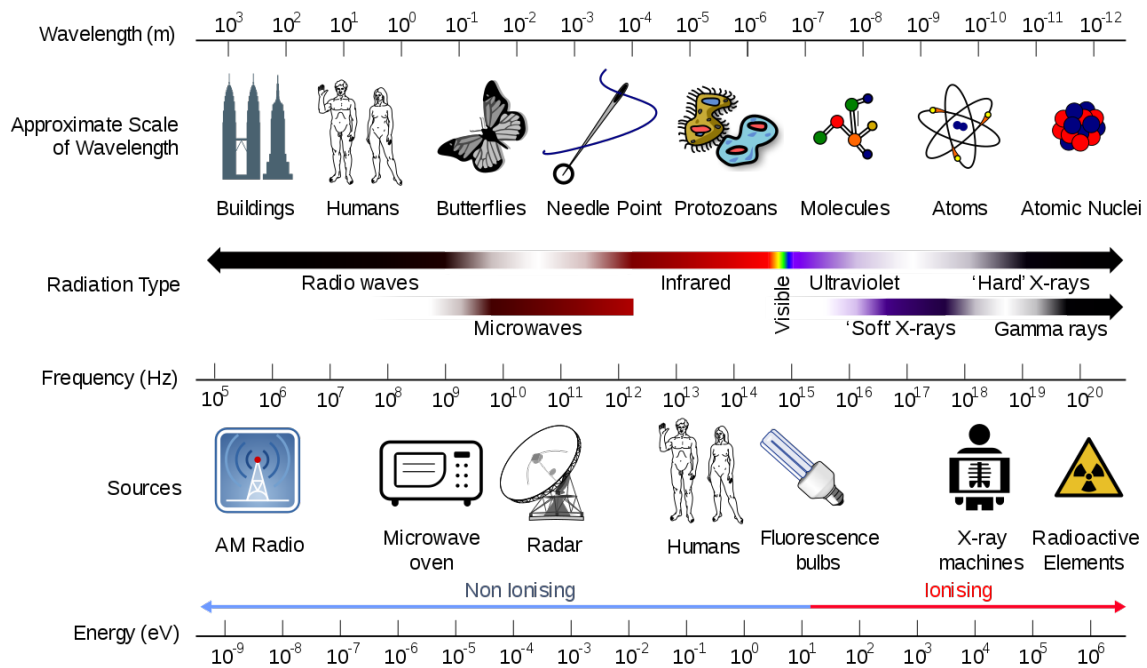


Figure 2-1: Overview of the electromagnetic spectrum showing energy per photon, frequency, wavelength, and their scale as well as the radiation types. Source: Reproduced from [85], under Creative Commons Attribution-ShareAlike License.

Methods using the electromagnetic spectrum are the most prominent in the field. Other energy sources, such as acoustic energy are used for thermoacoustic imaging [2], [24]. Ultrasonic energy for echocardiography [48], and electron energy used for electron microscopy [7], [51] are also of significant importance.

2-1-1. Microscopy

Microscopy is widely used in a variety of fields, from cell biology, through pharmaceuticals, to material inspections [29], [56], [63], [68], [77]. There are also numerous types of microscopy techniques. Examples of these can be seen in Figure 2-2 [68], [83], [87], [91].

The most common one is light microscopy [29]. Depending on the illumination and staining of the sample, one can find different specialized techniques, which mitigate certain issues and highlight specific sub-structures in the sample. The simplest approach is bright-field microscopy, seen in Figure 2-2a, where sample preparation is minimal compared to many of the more advanced methods. However, this simplicity comes at a cost, as problems with contrast and focus arise regularly. Neglecting special sample preparation, this technique magnifies the sample without highlighting or emphasizing any specific aspect or part of the sample, which is excellent for general, preliminary analysis. It can frequently provide points of interest for further analysis. A more advanced technique is dark-field microscopy [5] in Figure 2-2b, which provides significantly improved contrast in many cases, without staining of the sample. Another possibility is phase-contrast microscopy, seen in Figure 2-2c, where intracellular and extracellular structures become easily distinguishable due to the differences in their refractive indices.

Raising the complexity, samples are also often stained in order to highlight specific morphological structures as well as to increase the contrast between the structures. This can also boost the effectiveness of more straightforward methods. The most popular staining method uses hematoxylin and eosin (H&E), as seen in Figure 2-2d, where hematoxylin

stains cell nuclei purple, eosin colors the extracellular matrix and cytoplasm with a pink color [66]. It is important to note that there are cases where sample staining is either not possible due to biochemical limitations or infeasible due to increased preparation time and costs. In these cases, dark-field or phase-contrast microscopy may work well.

When separating these structures is not selective enough, more sophisticated methods are also available, such as fluorescence microscopy in Figure 2-2e. Among others, this permits staining with specific antibodies that are engineered to emit light with a particular wavelength, and this helps to identify unknown bacteria or the area of effect for a specific chemical species that the antibodies latch onto [57], [89]. Extensions to this are multi-staining approaches that simultaneously dye multiple compounds or structures.

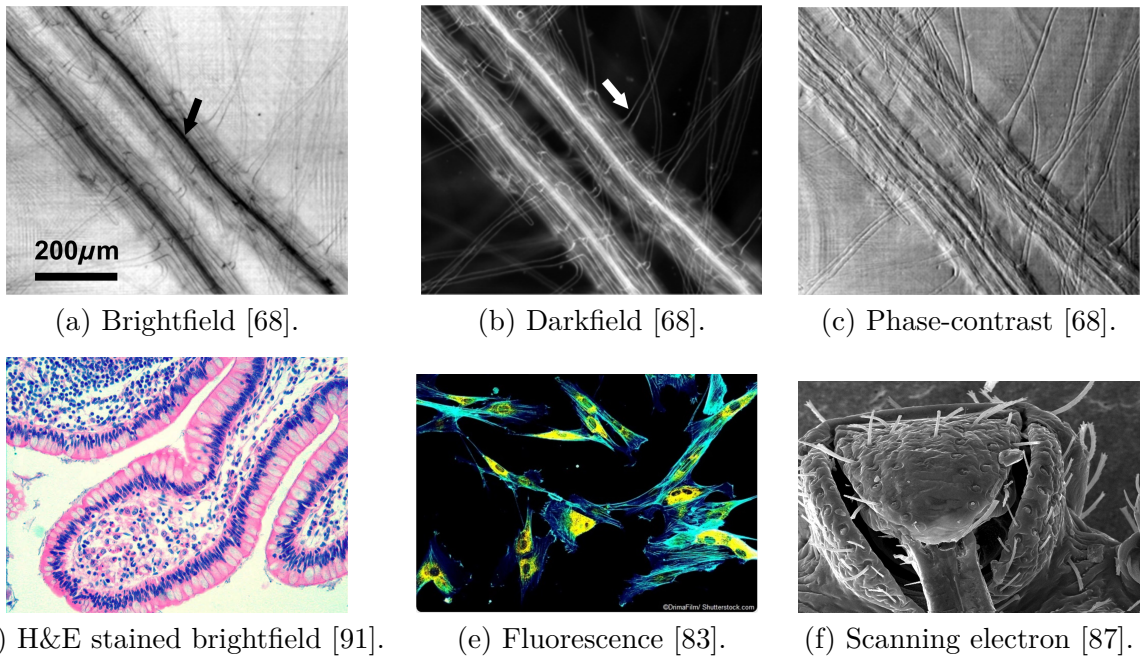


Figure 2-2: Examples of specialized techniques in microscopy. Source: [68], Figure 3, under Creative Commons Attribution License 3.0. Source: [91], under Creative Commons Attribution License 4.0. Source: [87], under Creative Commons License 1.0.

There are cases, however, where increased contrast and special staining is not enough, but raw magnification and high resolution are required [21], [51], [54]. When this happens, electron microscopy is a possible solution. These microscopes can achieve up to 1000-fold magnification compared to light microscopes. The transmission electron microscope (TEM) can be used to analyze viruses and microorganisms in the nanometer range. In contrast, the scanning electron microscope (SEM), visible in Figure 2-2f that allows for a three-dimensional view, but has lower resolution compared to the TEM.

A wide variety of noises can contaminate images made with these techniques. Some of these originate from physical properties of the measurement system, including but not limited to photon noise, which depends on the photon count, thermal noise, which is dependent on the temperature of the appliance and the integration time, and structured noise, which is affected by the differences in photon density inside and between viewports. Furthermore, other noise components come from the digitization process, such as readout noise, whose severity is sensitive to changes in the readout rate, and quantization noise, which is susceptible to the number of bits in the analog-to-digital-converter (ADC) [81], [94]. There are several other sources of noise in these systems as well. However, in this

document, our focus will be on the different types and some effective removal algorithms, rather than discussing the origins in-depth.

The sheer number of unique methods within microscopy can seem daunting, but this comes largely from the wide range of fields and applications they are used for, each with their own strengths and shortcomings.

2-1-2. Other imaging modalities

Even though microscopy is likely the most commonly used technique, it is almost always used in *in vitro*. In contrast, many other imaging techniques are capable of *in vivo* imaging, which is significantly less invasive and often a considerably quicker solution than going through surgery and sample preparation in order to acquire samples of interest.

One such technique is Magnetic Resonance Imaging (MRI), which uses short pulses of radio waves generated by a strong magnet [43]. When these pulses collide with different tissue types in the patient's body, a response wave is created. A computer is used to deconstruct the differences between the initial and response pulse, to determine their location. This makes it possible to create a 2-D tomographic image of the patient [81]. The principal sources of noise for this technique are radio frequency emissions in the patient's body due to thermal motion, instrument noise from the coils, and other electronics in the appliance [93].

Another important technique is Imaging Mass Spectrometry (IMS) [34], [47], which combines visual, microscopy-like imaging capabilities with detailed chemical information by measuring the mass-to-charge ratio of specific ions obtained from the sample. This technique can record the spatial distributions of a wide variety of atoms and molecules, including lipids, peptides, metabolites, and proteins. It also enables the visualisation of biomolecules throughout a tissue sample, e.g. protein distributions in tumor tissue, biomarkers [47]. In simple terms, the sample is ablated with electrons or photons from a laser, charging individual molecules. As a result, many of them fragment and release charged particles before returning to a more stable molecular state. These released particles can be captured as an ion signal that is then mapped as a function of their mass-to-charge ratio, resulting in a single mass spectrum per pixel. The measured mass spectra can be used to determine the chemical structure and composition of the sample, as well as for visualizing the spatial distribution of different chemical species across the sample [47].

Note that this is a simplified description of IMS, there are several different methodologies and instrumentations that may vary greatly, in terms of ion creation, mass selection, and detection compared to the description above. For example some methodologies are matrix-assisted laser desorption/ionization with time-of-flight mass analyzer, spark source mass spectrometry, thermal ionization-mass spectrometry, and accelerator mass spectrometry.

While this provides very detailed chemical information as well as spatial information on the sample, a high spatial and spectral resolution can create datasets in the range of hundreds to thousands of gigabytes for a single sample. Manual analysis of such results becomes cumbersome, if not impossible. For this reason, a considerable amount of work has been done to perform preprocessing, denoising, and analysis of these measurements in an automatic, data-driven manner both in the commercial and open-source communities. Examples of such tools are msIQuant [79] and Mass-Up [lopez-fernandez__mass-up: 2015](#).

2-2. Common noise types and removal

Noise can originate from several sources. Some of the most common ones are instrumental, environmental, and human-introduced [23], [36], [69]. Instrumental noise, for example, may come from the source of energy used for imaging, the spectral band captured by, and the quality of the imaging sensor. Others sources are environmental, such as ambient temperature of the room, temperature of the imaging sensor, and differences in background illumination between measurements. The final type of noise source is human-introduced, e.g. lack of proper maintenance and calibration, low-quality sample preparation, and general negligence regarding the proper use of the appliance and knowledge of the imaging technique. Many of these are introduced at a specific stage of imaging, most commonly during image acquisition due to light levels and sensor temperature, or during transmission, principally due to interference [81].

If users are not mindful of the necessary operating conditions of these imaging techniques, noise can make interpretation of the acquired images difficult, thus hindering the possibility of a correct diagnosis. These effects include but are not limited to unrealistic edges, vanishing or artificially added lines, corners, tiling, blurred objects, and disturbed background scenes [39], as seen in Figure 2-3. Understanding and modeling these common types gives us tools to characterize and invert them.

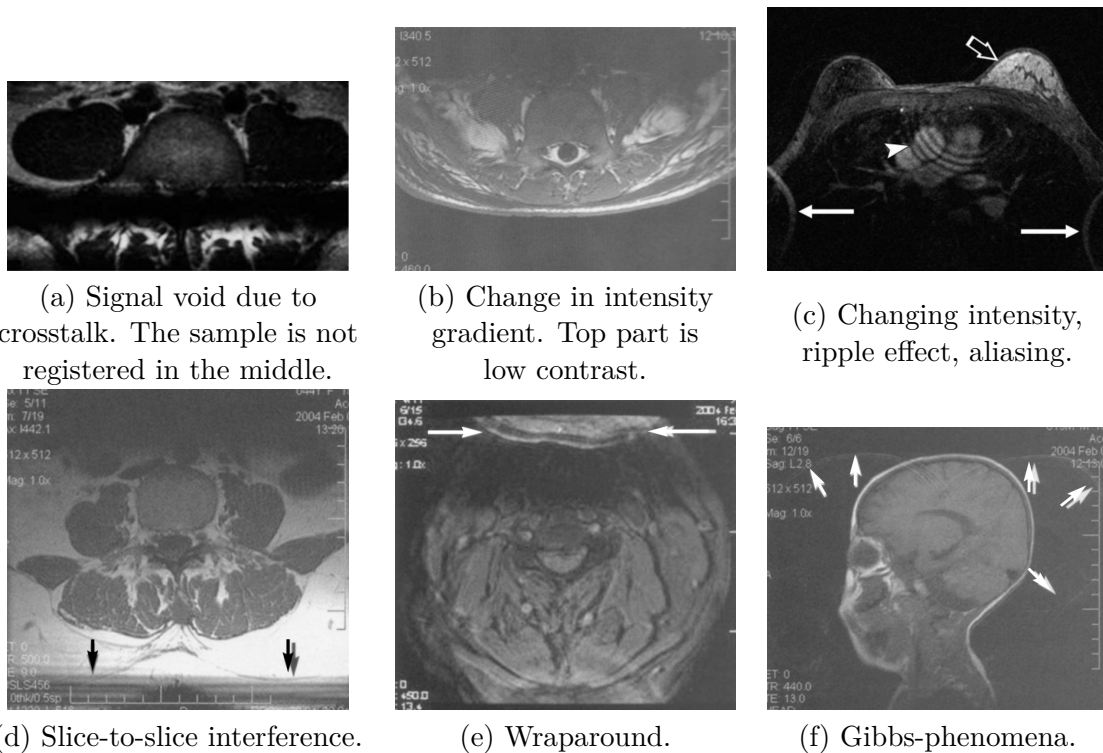


Figure 2-3: Examples of various noise artefacts on a variety of biomedical images [35], [59]. Source: [35], Figures 7,8,9,11, under Creative Commons Attribution License 4.0. Source: [59], Figure 17,25.

Pižurica et al. [45] argues that in the case of medical images, the aim of denoising is considerably different from other common domains and their applications, such as in entertainment or telecommunication. While in the latter, favorable characteristics of denoising are determined by viewing and hearing pleasure of users, for medical images, visually clean images may still disguise clinically vital features. Degradation of a crucial feature through

smoothing should be avoided at all costs, as well as the creation of artificial features that could lead to a faulty medical conclusion.

Requirements and parameters of the denoising also depend on the application. At the same time, a radiologist may find helpful information in the "texture" of a speckle-noise pattern upon manual examination [45]. A different type and level of denoising may be advantageous for an automatic segmentation algorithm, where performance heavily depends on the preservation and sharpness of edges and benefits from smoother textures.

Raw and corrupted images are often preferred by medical professionals, as obstruction of important features is always a possibility when the noise filters used are not sufficiently refined. At the same time, the same effect can be caused by the lack of filtering as well. Maximizing visibility and quality of features of interest is a delicate optimization problem, and the concerns above should be taken into consideration during the design of appropriate noise filters.

Even though other noise models, including but not limited to Rician, speckle, Rayleigh, gamma, uniform, and anisotropic noise are important, they are not discussed further in this chapter. More information on these can be found in literature, for example, in Gonzalez et al. [81].

2-2-1. Gaussian noise

The Gaussian noise model is frequently used in the domain, as it describes multiple sources of noise well and it is easy to describe and apply both spatially and in the frequency domain. The only variables in its distribution are the mean, μ , and standard deviation, σ , as seen in Equation 2-1 and in Figure 2-4b, which describes the corresponding probability density function (PDF) of x , a Gaussian random variable.

$$p(x) = \frac{1}{\sigma\sqrt{2\pi}}e^{-(x-\mu)^2/2\sigma^2} \quad (2-1)$$

Because of this convenience, it is often used in cases where it is only borderline applicable [81]. One such case comes from using the Gaussian model to describe noise generated by fluctuations in the light source or high temperature of the imaging sensor. This can cause noise to appear in the electronic circuit and sensor during scanning and digitization [81]. As a result of discretization and counting statistics involved in modeling noise in such electronic components, the resulting distribution can be considerably different from a Gaussian distribution. Even though the physical phenomena are modelled well by the Gaussian process, the digital, acquired signal is often not modeled accurately by the same model. In the practical sense, it shows up as intensity variation in the uniform regions of the image. For example this is seen during field emission SEM images, where the electron emission moves to another position, resulting in a fluctuation in beam current. This can often be minimized by choosing a scan time that is significantly shorter or longer than the fluctuation time of the source [80].



(a) Original Lena image.

(b) Lena image, with added Gaussian noise, with a mean of 0 and a variance of 0.01.

Figure 2-4: An example of the effect of Gaussian noise.

2-2-2. Impulse or salt-and-pepper noise

This type of noise is usually the result of impulse corruption, which is often significantly larger than the strength of the signal. In pixels of digital images, this results in randomly appearing, saturated, white or black pixels [81]. It is often caused by ADC errors as well as bit errors in transmission.

One possible description of impulse noise is shown in Equation 2-2, which describes the probability density function. For an 8-bit image $a = 0$ and $b = 255$ in most cases, while for a 16-bit image $a = 0$ and $b = 65535$, while P_a and P_b are their probability of occurrence. When P_a and P_b occurs, the pixel intensity is replaced by a and b , respectively. On a grayscale image, this resembles salt-and-pepper particles, which are randomly distributed over the image. An example of this type of noise can be seen in Figure 2-5b.

$$p(x) = \begin{cases} P_a, & \text{for } x = a \\ P_b, & \text{for } x = b \\ 0, & \text{otherwise} \end{cases} \quad (2-2)$$



Figure 2-5: An example of the effect of impulse noise.

This particular type of noise is extensively studied, and in the majority of cases, mean- and median-filtering is effective in removing it, thus enhancing the image quality [65]. The filter looks at the pixel and its immediate neighborhood, which is called a kernel, and takes the mean or median of the contained values, which then become the new pixel value in the filtered image. Small kernels are commonly preferred (3×3 or 5×5), as larger kernels smooth out the image and degrade edges as well as fine details. Median-filters, in particular, are significantly better at preserving edges. However, they may be less aggressive towards the noise, thus more likely to have residuals compared to mean-filters [53].

2-2-3. Periodic and quasiperiodic noise

Periodic functions are functions which repeat their values in set intervals. A formulation of this concept is shown in Equation 2-3, where $x(t)$ is a time-varying function, t is time, and P is the period of time in which the repetition occurs. An alternative formulation is shown in Equation 2-4. For this type of noise, the frequency domain often gives a concise description of the signal, in the continuous case, a sine or cosine function with a single frequency appears as a peak in the frequency domain, assuming a one-sided transform.

$$x(t) = x(t + P) \quad \forall t \in \mathbb{R} \quad (2-3)$$

$$|x(t) - x(t + P)| = 0 \quad \forall t \in \mathbb{R} \quad (2-4)$$

The theory of almost periodic functions was created by Harald Bohr, it has been extended and further developed by multiple authors, such as Vyacheslav Stepanov, Hermann Weyl, Abram Samoilovitch Besicovitch, and John von Neumann, among others [9], [92]. The definition of quasiperiodicity used in this thesis is closest to Bohr's definition of uniform almost periodic functions. However, existence and uniqueness theorems in particular spaces, such as the Banach space, and general properties of the solution are not discussed,

as it is not the primary focus of this thesis, further information on these topics can be found in [9]. In the signal processing sense, the concept can be explained relatively easily, but the mathematical description and its consequences are somewhat more complicated. In principle, an quasiperiodic signal is locally virtually periodic, meaning that each period is identical to its neighbors. However, globally, the similarity of such periods can change. Looking at the frequency domain representation, this results in a distribution around some peak. A mathematical formulation of this phenomena is described in Equation 2-5 and an alternative formulation in Equation 2-6. This formulation is similar to Equation 2-3 and 2-4, except that the period is not exact, which is represented by ϵ , the error term. This definition originates from Bohr's definition of uniform almost periodic functions, however, intentionally less precise. It serves as a conceptual description that imposes no theoretical restrictions on the error term ϵ , in order to incorporate structured noise, such as viewport-stitching, discussed in Section 2-2-4.

$$x(t) \approx x(t + P) \quad \forall t \in \mathbb{R} \quad (2-5)$$

$$|x(t) - x(t + P)| = \epsilon \quad \forall t \in \mathbb{R} \quad (2-6)$$

In the practical sense periodic and quasiperiodic noise often appears as an additive repeating pattern over the image, as seen in Figure 2-6b, which appears in the form of repetitive crests and valleys. When looking at the frequency domain representation, crisp or distributed peaks.

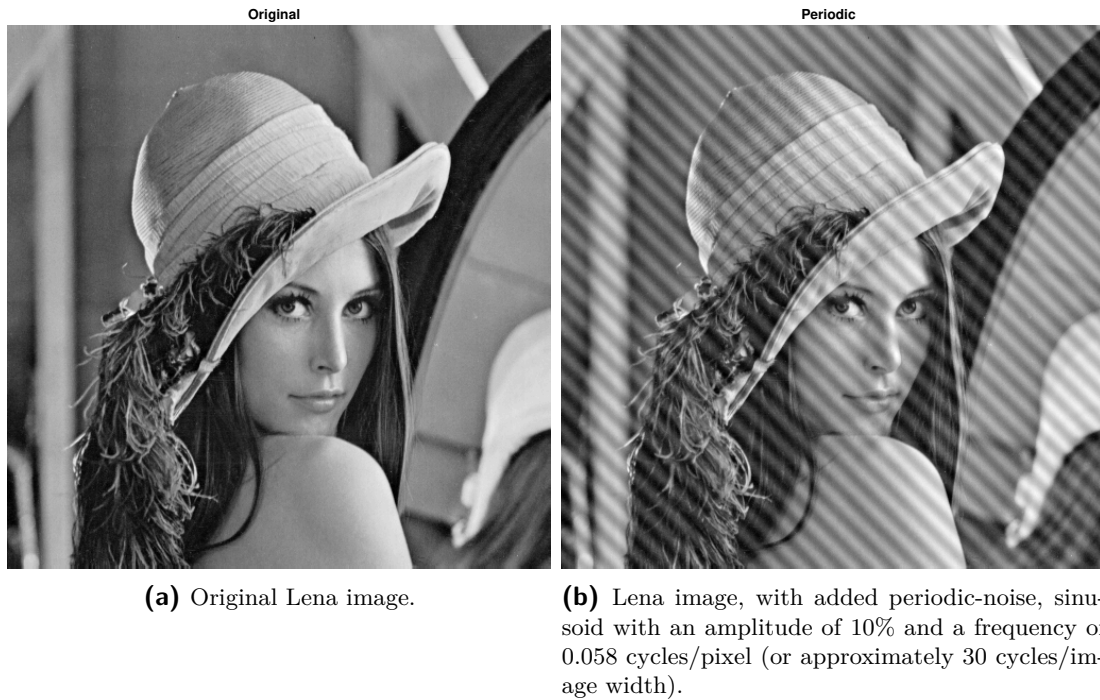


Figure 2-6: An example of the effect of periodic noise.

The source of this noise is often electrical and electro-mechanical interference [81]. For a specific sensor or appliance, expected parameters and the probability density function of this noise may be given in the characterization sheet or data sheet of the sensor. However, estimation of the parameters is also possible in case the imaging can be done on a uniform, flat sample, which provides a good baseline for the method noise of the system [32].

An interesting sub-phenomenon is the Moiré effect, which is an interference pattern that occurs when multiple periodic or quasiperiodic structures are overlaid, thus creating a new superimposed structure. An example of this can be seen in Figure 2-7. These patterns are frequently also seen as aliasing phenomena [32].

Filtering of such corruption can be done in different domains, such as spatial [32], frequency [70], [72], [74], [82], and the wavelet domain [84]. A detailed description and explanation of the fundamentals of each can be found in Chapter 2-3.

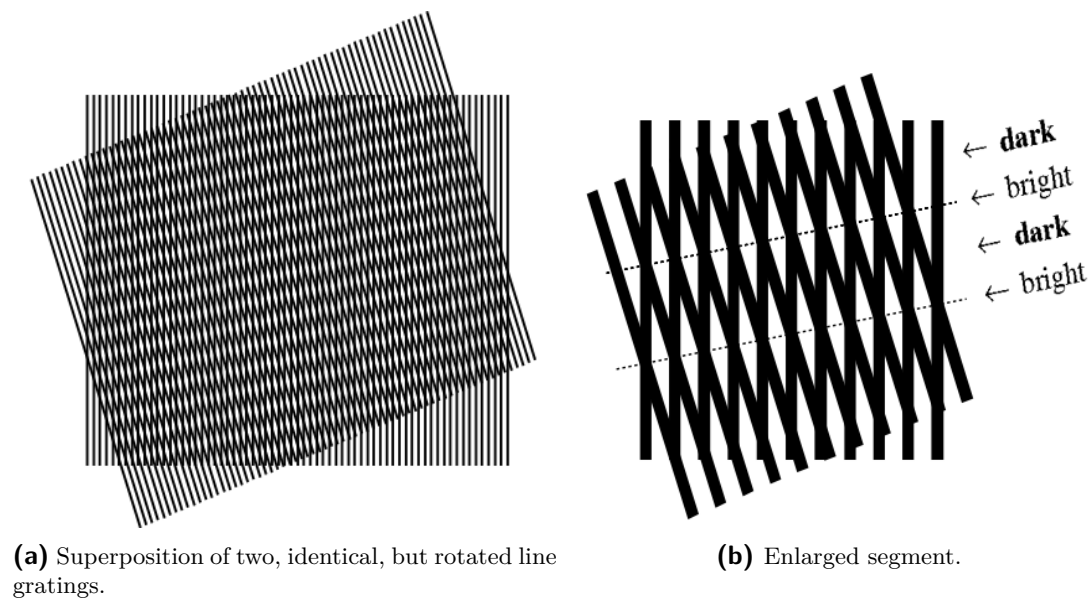


Figure 2-7: Examples of Moiré patterns. Source: [86], Figure 1.1, reproduced with permission from Springer-Verlag London Limited.

2-2-4. Structured noise

An instance of structured noise is a phenomenon called viewport-stitching. High-resolution microscopy images, for example, are not often acquired by a single high-resolution imaging sensor, at a single instant in time. They are often acquired by a lower-resolution sensor that takes an image of a small part of the sample, this is what we refer to as a viewport. This sensor is then moved along a specific structured path, such that a lower-resolution image is taken of each and every part of the sample. When all the areas of the sample have been captured, the lower-resolution images or viewports are connected or stitched together, this results in a high-resolution image. The amount of light that falls on the sample within and across viewports is often not uniform, especially without regular calibration and maintenance. When stitching these non-uniformly lit viewports together, a structured, repetitive pattern emerges, which is referred to as viewport-stitching artifact or viewport corruption [27].

As seen in Figure 2-8a, this results in a repeating pattern emerging between the adjacent viewports and a gradual vanishing in certain directions [56]. The differences between Figure 2-8a and 2-8b include a more uniformly distributed illumination within the viewports, a considerable change in the size of viewports as well as post-acquisition color enhancement of the staining on 2-8b. It is important to note that differences across viewports are still present even after enhancement.

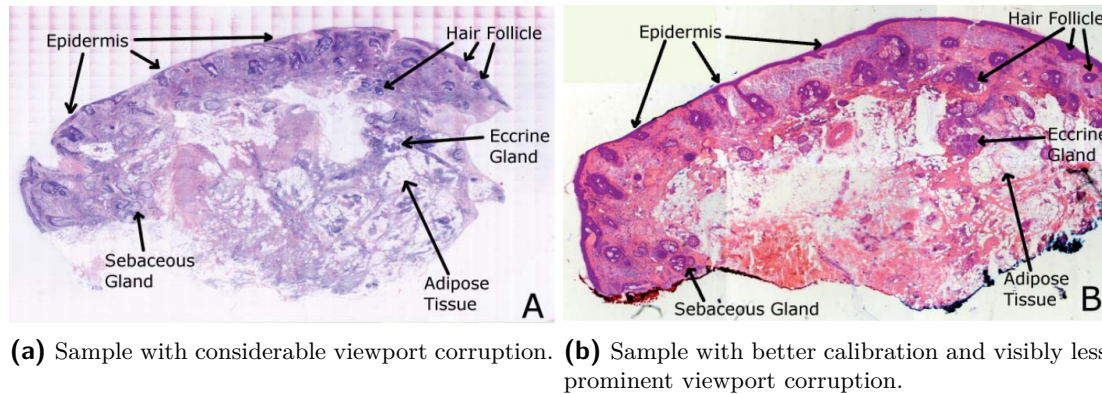


Figure 2-8: Viewport-stitching artifact on H&E stained microscopy images. Source: [56], Figure 5, under Creative Commons Attribution License 4.0.

2-3. Decomposition and domain transformation

In a general denoising pipeline, common steps include decomposition, noise component identification, and filtering, followed by reconstruction. Decomposition commonly happens either through a domain transformation such as Fourier or wavelet transform or through factorization methods. The choice of decomposition also often dictates the type of reconstruction or inverse transformation that is required. The choice of decomposition, identification, and filtering methods are strongly tied to the goals of the denoising and the specific characteristics of the noise.

Periodic and quasiperiodic noise, as introduced at the end of Section 2-2-3, can be approached through a wide range of methods, both in terms of domain transformation and factorization. In this chapter, exploration and comparison of these methods takes place with a focus on their strengths and weaknesses with regards to decomposing corrupted medical images.

2-3-1. Signal domains

Using transformations, signals can be represented in a different basis than the one in which they were acquired in. For example, a real-valued time-domain signal reports how its measured variable evolves as time progresses. Using a Fourier transform, we can obtain a complex-valued frequency domain representation of that same signal. This representation shows the frequencies, amplitudes, and phases of the sine and cosine waves that, when summed, build up the original signal.

There are a large number of possible transformations available, and this section expands on the most popular and more effective transforms that can capture periodic and quasiperiodic signals specifically. These methods can be categorized into spatial, Fourier, and wavelet domain models.

2-3-1-1. Spatial-domain parametric models

In relatively simple, clearly periodic cases, it is possible to use the Fourier series expansion and estimate the coefficients of the basis functions using flat, low variance patches of the image, directly in the spatial domain. Such patches can be for example, a uniform patch of grass in a photo of a football match, or a viewport outside the biological sample in a microscopy image. In such areas, assuming that the mean of the flat, low variance patch approximates the true, noiseless image well, the overarching pattern can be reduced to

a system of linear algebraic equations [32]. This method can be expanded and applied to the whole image. Even though this sounds straightforward and effective, in practice, the useability of this approach is quite limited. Images often have no patches where the mean approximates the noiseless image well, or the coefficients found within the patch may not apply uniformly to all other parts of the image, which is a common observation in quasiperiodic noise cases. Finally, the accuracy of such fitting is sensitive to the noise that is present. As mentioned in the conclusion of Sidorov et al., where both the spatial domain and a Fourier-based method was used, the latter tends to provide a higher quality removal. It is important to mention that in [32] the goal was denoising in a way that is pleasing for the human eye from a broadcasting perspective, not from the standpoint of clinical usability, as discussed in Section 2-2.

2-3-1-2. Fourier-transform and frequency-domain analysis

Jean Baptiste Fourier published his paper on heat propagation in solid bodies in 1907 [1]. He discovered that every continuous or discontinuous function can be composed using an infinite number of sines and cosines. The continuous Fourier-transform (CFT), however, has a significant challenge in practice, due to the integrals, some of which cannot be represented in a closed form, making the analytical solution unattainable in some instances. The discrete Fourier-transform (DFT), which resolves this issue and thus is more readily applicable in many practical cases, also begun its development around the same time [64].

The DTF has been applied to a wide variety of problems, in virtually every scientific and technological field, including but not limited to optics, imaging, electricity, acoustics, communications, signal processing, biological engineering, hydrodynamics, heat propagation, mechanics, geophysics, spectroscopy, statistics, and mathematics [64]. Even after the widespread availability of digital computers, the computational costs of the algorithm were still a burden. This limitation was mitigated significantly by the introduction of the fast Fourier-transform (FFT) in 1965 [8], an efficient algorithm for calculating the DFT.

Effective analysis and filtering of periodicity can typically be done more easily in the frequency domain, which in terms of the Fourier transform means that periodic signals can be deconstructed to a combination of specific sines and cosines. The Fourier-domain coefficient locations are dictated by the frequencies, and their amplitude is proportional to the magnitude and phase of these deconstructed trigonometric functions. A common approach for fairly well-localized noise is to use a filter to capture it. The simplest of these filters are band-reject, band-pass, and notch filters [81]. Such filters and other possibilities are detailed in Section 2-4-2.

The 2-D discrete Fourier transform implementation we used is shown in Equation 2-7. For an $n \times m$ matrix in the spatial domain, X , and its discrete Fourier transform Y , $\omega_m = e^{-\frac{2\pi i}{m}}$ and $\omega_n = e^{-\frac{2\pi i}{n}}$ are complex roots of unity, where i is the imaginary unit, p and q are indices in the Fourier domain, j and k are indices in the spatial domain, p and j are indices that run from 0 to $m-1$, and q and k are indices that run from 0 to $n-1$.

$$Y_{p,r} = \sum_{j=0}^{m-1} \sum_{k=0}^{n-1} \omega_m^{jp} \omega_n^{kq} X_{j,k} \quad (2-7)$$

As the DFT is a discrete and finite numerical approximation of its continuous counterpart, some disparity is inherent. The most critical phenomena that arise from the discrete and finite numerical approximation are sampling introduced ambiguity and the necessity of finite windowing, which are distinct characteristics of the DFT. Both of these cause a

certain amount of information to be lost, consequentially bringing their own artifacts [64].

Sampling introduced ambiguity causes aliasing, which means that continuous signals with vastly different frequencies can deliver precisely the same discrete signal after they are sampled. An example of this is shown in Figure 2-9. In order to counteract this, the frequencies of interest have to be at most half the sampling frequency. Above this limit, the signal has to be band limited, meaning that the frequency components have to be zero. This is also known as the Nyquist-Shannon sampling theorem [12]. Anti-aliasing filters can mitigate the effects of improperly filtered and sampled signals. More information on the topic is provided by Gonzales et al. [81]. To completely avoid this ambiguity, a continuous signal has to be sampled frequently enough such that the captured digital waveform is unique. In practice, the sampling frequency has to be more than twice as large as the highest frequency of interest. As a consequence, the higher-frequency signal shown on Figure 2-9 is sampled insufficiently, as the discrete waveform is not unique, while the lower-frequency signal is. A broader analysis and description of the phenomenon is presented by Amidror in Chapter 5 [64].




Figure 2-9: The sampled version of two different signals is identical. The higher-frequency signal shown on Figure 2-9 is sampled insufficiently, as the discrete waveform is not unique, while the lower-frequency signal is. Source: [64], Figure 5.1, reproduced with permission from Springer-Verlag London Limited.

While the CFT is applied with an infinite range, the DFT is only able to look at a finite window within the original range and thus assumes a spatially cyclic behavior between the first and last elements of this range. This finite range can be seen as a signal domain multiplication of the signal of interest with a rectangular windowing or apodization function. A windowing function can have different shapes, the simplest ones are rectangular or triangular, but many of them follow some variation to a bell shape. Finite windowing causes leakage, which in principle means that the Fourier-spectrum of a continuous signal can be significantly different from its discrete counterpart. This phenomena is illustrated on Figure 2-10, though an example described in the next paragraph. In the frequency domain this multiplication causes side-lobes to appear, which is already an imprecise representation of the continuous signal. In the signal domain, this multiplication can cause discontinuities between the first and last elements of the function, which in terms of the

frequency domain, results in leakage. This causes representations of precise peaks to be distributed to a number of adjacent peaks. In Chapter 6 of the book by Amidror [64], this phenomenon is described in more detail.

To illustrate leakage, an example is shown in Figure 2-11. The DFT of $g(x) = \cos(2\pi f x)$, where $f = \frac{18}{32}$ is sampled at 4 Hz, with a rectangular apodization window $w(x) = \text{rect}(\frac{x}{R})$. Here, x is the signal domain input and $R = 32$, which is the sampling range length or sample count as seen on the top axis of Figure 2-11a and the $\text{rect}()$ function is defined as in Equation 2-8. Next, the windowed $g(x)w(x)$, is fed through the DFT, resulting in $G(u) * W(u)$. Here u is the spectral domain input, $*$ is the convolution operator, $g(x)$ and $G(u)$ are the signal domain and spectral domain representations of the same function. The transformation from $g(x)w(x)$ to $G(u) * W(u)$ has two significant consequences. First, because of the rectangular window, the impulses of the original signal are changed, this is shown on Figure 2-10b and 2-10c, this corresponds to the shape of $W(u) = R\text{sinc}(Ru)$, where the $\text{sinc}()$ function is defined as in Equation 2-9. This causes the emergence of previously non-existent side-lobes or ripples, the comparison is shown in Figure 2-10a and 2-10c. Secondly, depending on the sampling frequency, which determines the frequency resolution, the maximum of the impulse is only captured in case the function frequency f is an exact multiple of the frequency step. This is not the case for Figure 2-11a and 2-11b, as $f = \frac{18}{32}$ is 4.5 frequency step for every full cycle of $g(x)$, which is not an exact multiple. In most cases, several, smeared values appear around the continuous peak, as seen in Figure 2-11b [64].

$$\text{rect}(x) = \begin{cases} 1, & -\frac{1}{2} < x < \frac{1}{2} \\ \frac{1}{2}, & x = \pm\frac{1}{2} \\ 0, & \text{otherwise} \end{cases} \quad (2-8)$$

$$\text{sinc}(x) = \begin{cases} \frac{\sin(\pi x)}{\pi x}, & x \neq 0 \\ 1, & x = 0 \end{cases} \quad (2-9)$$




Figure 2-10: Examples of the effect of finite windowing both in the signal and frequency domain. Source: [64], Figure 6.2, reproduced with permission from Springer-Verlag London Limited.

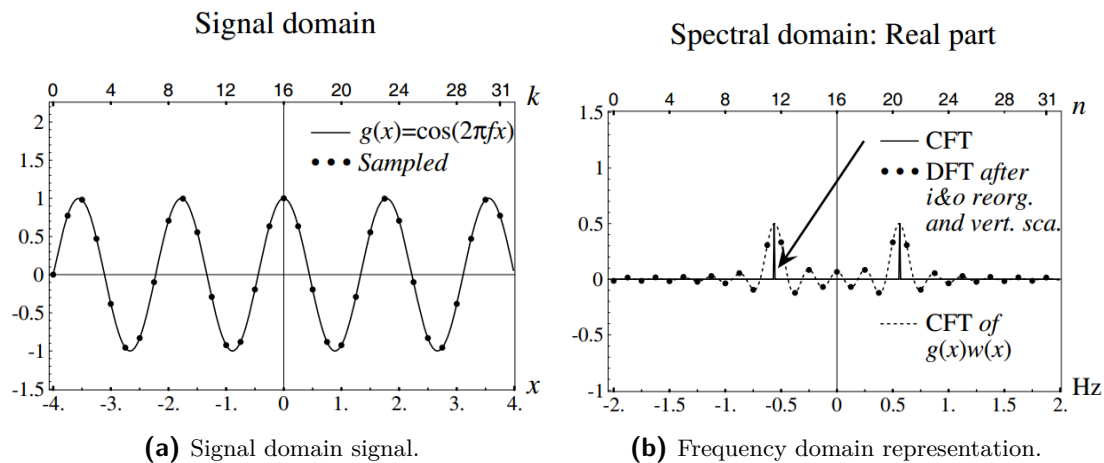


Figure 2-11: Examples of leakage both in the signal and frequency domain. Source: [64], Figure 6.3, reproduced with permission from Springer-Verlag London Limited.

2-3-1-3. Wavelet-transform

The wavelet transform follows a similar principle as the Fourier transform, as it breaks down arbitrarily complex signals to simpler ones, although in a significantly different way

[18]. While the Fourier transform decomposes signals to a multitude of sine and cosine functions that are infinite by nature, the wavelet transform uses waves with a finite range, which rise from 0, then quickly decay to 0. The difference in representation between the wavelet and Fourier transform is significant, as the wavelet transform gives a considerably more temporal oriented description, which is better applicable to capturing behaviors of time-varying, non-stationary signals.

The wavelet-transform deconstructs the signal into basis functions, which are expanded, shrunk, and shifted versions of a single function, called the mother wavelet. Even though there are no frequencies in this transform, scales capture a similar phenomenon. It describes, in what way the mother wavelet expanded, shrunk, and shifted to represent the signal. This gives rise to various scales with various levels of resolution [60]. The choice of the wavelet entails that decomposition of the same signal, with different mother-wavelets, results in different representations. Since the wavelet-transform is not used in QID, the exact formulation of the transform is omitted. The introduction of the concept, however, is important as it is used by a number of denoising algorithms.

It is less popular and common than the Fourier-transform. The wavelet-transform is commonly used in areas such as signal analysis, image processing, compression, and denoising [18], [55]. Pižurica et al. [45] reviewed previously published applications for wavelet denoising of MRI and ultrasound images of brains, explicitly focusing on speckle-noise both from the perspective of signal processing and clinical usability, as discussed in Chapter 2-1. They concluded that wavelet domain denoising could improve the human readability of MRI images, as well as the automatic segmentation of images. However, a different set of parameters is used for each application.

Ionita and Coanda [84] proposed a method specifically for quasiperiodic corruption removal in microscopy images, using the wavelet-transform as an extension to their previously proposed method using the Fourier-transform [70], [82]. Their findings show that the wavelets handle noise-related non-stationary behavior better, and overall provide slightly better results in terms of peak signal-to-noise ratio (PSNR). A detailed comparison of tuning sensitivity, differences in computation times, and quantization of non-stationarity was not included.

2-3-2. Matrix factorization methods

Matrix factorization or decomposition can be, in simple terms, the writing of a matrix as a product of matrices with specific characteristics. For example, lower-upper decomposition (LUD) separates a matrix into a lower and an upper triangular matrix [19]. Another example would be the eigendecomposition of a matrix, which creates linearly independent eigenvectors and their corresponding eigenvalues that then can be used e.g. to calculate the inverse of that matrix.

In most cases, matrix decomposition methods capture a specific pattern or structure in the data of the matrix, which then can be used for further computation and optimization. In this chapter, a few of the more popular factorization methods and their applicability to periodic noise separation are discussed.

2-3-2-1. Principal component analysis

Principal component analysis (PCA) is a technique that provides a basis for many multivariate data analysis applications. Applications range from dimensionality reduction, variable selection and classification, to prediction, decomposition, and outlier detection [13]. When the data is multivariate and the original variables are interrelated or corre-

lated, finding the most impactful features and combinations that capture and summarize the relationships between the original variables can be a complicated task. In order to achieve this, PCA orthogonally transforms the basis of the original dataset to a new, uncorrelated set of linear combinations of the original feature space, these are called principal components (PC) [78]. The PCs are ordered based on the total variance explained throughout the whole dataset, which means that the original variables that are combined together within a PC are highly correlated with each other. The first few principal components capture the majority of all the variation present in the data. This is illustrated along two dimensions in Figure 2-12. Internally, PCA is formulated as an eigenvalue problem and the new set of basis vectors are the eigenvectors of the covariance matrix. Ordering them in decreasing order of the corresponding eigenvalues, it ensures that the PCs successively capture directions with the maximum remaining variance of the data [31].

Consider X an $n \times m$ data matrix, it contains n observations and m dimensions. W is an $m \times m$ weight matrix, its columns are the eigenvectors of $X^T X$, such as w_1 with size $1 \times m$. Y is an $n \times m$ matrix containing the principal component scores.

$$Y = XW \quad (2-10)$$

When the algorithm is used for dimensionality reduction, only the first s eigenvectors in W are retained, resulting in $W_s = [w_1, \dots, w_s]$. W_s is an $m \times s$ matrix, effectively using only the first s principal components, as seen in Equation 2-11. This results in Y_s , which is an $n \times s$ matrix, often significantly reduced in size compared to Y , while still explaining a significant portion of the variance. Choosing the number of PCs to retain is specific to the task at hand, but commonly it is determined by setting a threshold in terms of variance explained. For example, one tries to retain the smallest number of PCs that still explain at least a set percentage of the variance for a specific dataset, commonly e.g. 90% or 95%.

$$T_s = XW_s \quad (2-11)$$

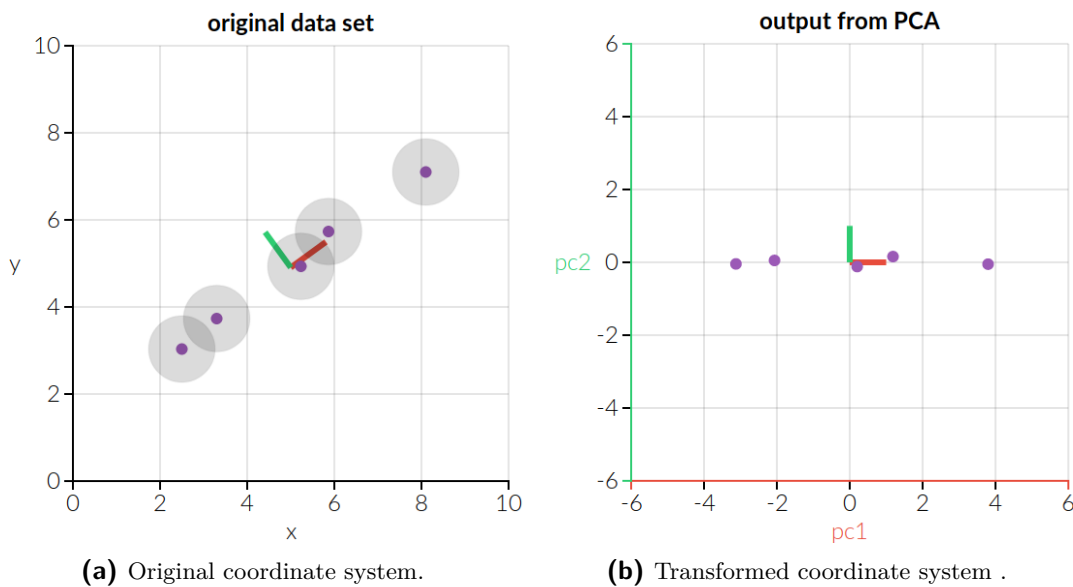


Figure 2-12: Two-dimensional example of coordinate system transformation using PCA. PC1 captures the majority of variance in the dataset. While PC2 still captures some variance, it is considerably less, compared to PC1 [90].

2-3-2-2. Non-negative matrix factorization

Non-negative matrix factorization (NNMF) is a method where all matrices in the decomposition are constrained to be non-negative. Its area of application is similar to PCA, as it can be used for dimensionality reduction, decomposition, and outlier detection [26]. However, the mode of operation is fundamentally different, compared to PCA, which learns a more holistic representation while NNMF provides a strictly additive, parts-based representation of the data. When used on a dataset of facial images, this results e.g. in each factor being a specific facial feature, such as eyes, eyebrows, mustache, a shadow on either side of the nose [25].

In terms of image denoising, this parts-based representation can lead to some combination of the factors capturing purely noise-related components.

2-3-2-3. Robust principal component analysis

Robust principal component analysis (RPCA) is an extension or modification of PCA [78]. It assumes that the $n \times m$ data matrix, M , can be decomposed as a superposition of a low-rank matrix, L , and a sparse component, S . In a real-world use-case, L may represent a face in a facial recognition task, or the background in video surveillance, while S may correspond to noise, such as a shadow on the face in a facial recognition, or a person moving through the footage in video surveillance [78]. As PCA is sensitive only to variance, thus it is also sensitive to outliers. An individual, heavily corrupted observation in the dataset may render the estimation arbitrarily far from the true low-dimensional representation [58].

RPCA can be formulated as a convex minimization problem in the form of Equation 2-12, a surprising discovery by Candes et al. [58] is that a universally correct tuning parameter can be used, which can exactly recover the low-rank and sparse component with high probability, the parameter setting is shown in Equation 2-13. The algorithm's applications include video surveillance, face recognition, latent semantic indexing, as well as ranking and collaborative filtering [58]. This same concept may be applicable to the separation of quasiperiodic and structured noise, from the underlying signal when used on the Fourier spectrum of an image.

$$\begin{aligned} \min \quad & \|L\|_* + \lambda \|S\|_1 \\ \text{subject to} \quad & L + S = M \end{aligned} \tag{2-12}$$

$$\lambda = \frac{1}{\sqrt{\max(n, m)}} \tag{2-13}$$

2-4. Filtering and identification

A general denoising pipeline often includes decomposition, noise component identification, and filtering of the noise components, followed by reconstruction.

Filtering is strongly connected to the characteristics of the noise to be removed and is highly application-specific, unlike the decomposition and factorization methods detailed earlier. The effectiveness of identification and filtering is heavily dependent on the assumptions held about the clean signal or image, as well as the characteristics and distribution of the corruption.

2-4-1. Recognition and clustering of outliers

Once the transformational and compositional steps are conducted, a decision has to be made on what part of the measurement to characterize as noise and what is expected to be the remaining signal of interest. This is an application-specific characterization, and the decisions made are based mainly on assumptions about the characteristics of noise and true signal.

Some tools available for recognition and labeling of outliers are explored in this section. Statistical analysis has to deal with the outlier detection problem quite frequently, thus some of its methods offer a broad toolkit for handling this challenge, further complemented by thresholding methods designed for applications in specific domains.

In statistical outlier detection, a priori knowledge of the distribution of a variable or data set often determines the range of tools that can be used. Most real data sets are corrupted by errors and the assumed distribution may not describe the data set well. It is crucial to determine if a measurement that is significantly different from the expected value comes from an error or from the expected, normal variability inherent to the data. Significantly different measurements are called outliers and naively including them in statistical analysis can heavily influence the results, often for the worse [42]. To counteract this effect, there are ways to detect and handle such observations. Some of the most basic and most commonly used techniques are considered here.

2-4-1-1. Z-score and three-sigma rule

In simple terms, a normal distribution is characterized by its mean, μ , and its standard deviation, σ . Looking at an observation x_i , one can calculate the ratio z_i as in Equation 2-14. The ratio z_i is the z-score that is the distance between the observation and the mean, divided by the standard deviation of the distribution. The likelihood of $|z_i| > 3$, assuming a normal distribution, is approximately 0.3%. This means that the absolute difference of the observation to the mean exceeding three standard deviations is very unlikely, however, not impossible. Accordingly, the "three-sigma rule" is a quick rule-of-thumb that suggests when observations could be considered outliers [42].

$$z_i = \frac{x_i - \mu}{\sigma} \quad (2-14)$$

In more detail, after defining a range around the mean, the probability of a sample falling into that range can be derived based on the cumulative distribution function of the Gauss-Laplace normal distribution, as shown by Grafarend [41]. The commonly used result of these calculations is that in a normal distribution approximately 68% of the samples fall within one standard deviation from the mean on either side, 95% within two and 99.7% fall within three standard deviations.

This approach works reasonably well in practice by substituting the distribution mean and distribution standard deviation with the sample mean and sample standard deviation. However, some considerations and caveats should be recognized. Even though the likelihood of a "three-sigma outlier" is low, it does not mean that it is impossible or that they are incorrect observations. Large samples should contain some genuine yet unlikely observations. Q-Q plots may be better suited for analysis of unexpected behavior in the distribution. When using the three-sigma rule heuristic on a small sample, the mean and standard deviation of the sample are not sufficient descriptors of the assumed distribution. For example, in a sample of 10 observations, $|z_i| < 3$ for every observation, independent of the value of the observations [42].

The z-score is widely used for hypothesis testing in empirical sciences. The significance of observations are directly dependent on the chosen confidence level. In order to qualify a hypothesis based on a set of observations as discovered, the confidence level required changes in different fields of application. Within social sciences, it is commonly 95% or two-sigma, while in particle physics, it is 99.99994% or five-sigma. It is important to note that the chosen confidence level is one of the parameters that determine the required sample size of the population. Furthermore, true outliers affect both the sample mean and the sample standard deviation, and this may make the method somewhat unreliable in cases where the normal distribution assumption is violated [42].

2-4-1-2. Modified Z-score or robust three-sigma rule

Outliers can significantly affect the sample mean and standard deviation, thus the sampled and theoretical measures of a distribution can be significantly different. In order to have more robust outlier detection, robust measures have to be used.

In the case when the normality assumption is held, but sensitivity to outliers should be decreased, a simple, robust modification to the three-sigma rule is available. Instead of using the mean, the median \tilde{x} is used and for a dispersion measure, instead of the standard deviation, the normalized median absolute deviation (MADN) is used, as shown in Equation 2-15 [15]. This modification describes t_i , as shown in Equation 2-16 and outliers are identified with threshold $|t_i| > 3$. However, Iglewicz et al. [14] suggests a threshold of 3.5. In this normalized case, the outlier detection behaves the same way for the theoretical normal distribution. In practical cases, it behaves very similarly to the three sigma approach, except it is less sensitive to outliers [42].

$$MADN(x) = \frac{\text{median}(|x_i - \tilde{x}|)}{0.675} \quad \forall i \in \{1, \dots, n\} \quad (2-15)$$

$$t_i = \frac{x_i - \tilde{x}}{MADN(x)} \quad \forall i \in \{1, \dots, n\} \quad (2-16)$$

2-4-1-3. Interquartile range

Another robust measure uses the quartiles of the sample. The first quartile of a sample is defined as the median of the set that ranges from the smallest number in the set to the median of the set. The third quartile is the median of the set that contains everything from the median to the largest number. The interquartile range (IQR) is the difference between the third and first quartile. In other words, the range where the central 50% of the data is located [42].

For outlier detection using box plots, subtracting 1.5 IQR from the first quartile and adding 1.5 IQR to the third quartile determines the normal range of values. Any observations lower or higher than that, respectively, are considered to be outliers. Since the quartiles and the IQR are not sensitive to occasional, large or small values, this method is robust to outliers. For a theoretical normal distribution, this range covers approximately 99.3% of the area.

It is important to note that this method does not assume normality and can be used for skewed distributions as well. The first and third quartiles are not symmetrical around the mean or median, unlike in the case of the aforementioned methods, but are strictly dependent on the sample distribution [42].

2-4-1-4. Hard and soft thresholding

With domain transformations such as the Fourier or the wavelet transform, certain types of noise and the signal of interest behave distinctly. With the Fourier transform applied, periodic noise tends to be concentrated in a handful of coefficients, while the true signal is usually broadly distributed throughout the spectrum. This is different in the wavelet-transform, where the signal of interest is usually more concentrated in specific bands, while noise is distributed among multiple bands [40].

To exploit this difference in behaviour in either case, thresholding approaches have been developed to classify or label specific transform coefficient, components, and areas as signal or noise using a specific set of rules. This can lead to the development of targeted filtering approaches to remove the unwanted patterns while preserving the useful ones. More sophisticated methods have been proposed to mitigate some of the caveats when using the wavelet transform for noise filtering. However, hard and soft thresholding are still one of the most effective, prevalent, and reliable techniques available [60].

According to Legendijk et al. [40], when a signal is decomposed adequately into spatial and temporal bands with different resolutions and orientations, the noise components are spread out over the bands. In contrast, the true signal is concentrated in specific, individual bands. Truncating components under a threshold to zero, while retaining the ones over said threshold may eliminate the noise with minor degradation to the true signal.

The main difference between hard and soft thresholding functions, also called coring functions, seen in Equations 2-17 and 2-18 respectively, is that the former is discontinuous while the latter is not, as shown in Figures 2-13a and 2-13b. In terms of visual artifacts, the soft thresholded signal or image generally appears to be smoother. However, this may not be beneficial for all applications, specifically medical imaging, as it may hide important details. In the mathematical sense, soft coring also entails better statistical properties compared to its hard counterpart, detailed by Donoho et al. [16], [17]. In these publications, they derive SureShrink, which is a soft thresholding based smoothness adaptive noise suppression algorithm. Some of its advantages include continuation of the rule and simplicity of the Stein's Unbiased Risk Estimate (SURE). Ergen [60] noted that hard thresholding is more sensitive to small changes in the signal, thus denoising performance may be less stable.

$$y(n) = \begin{cases} x(n), & |x(n)| \geq T \\ 0, & |x(n)| < T \end{cases} \quad (2-17)$$

$$y(n) = \begin{cases} \text{sign}(x(n))(|x(n)| - T), & |x(n)| \geq T \\ 0, & |x(n)| < T \end{cases} \quad (2-18)$$

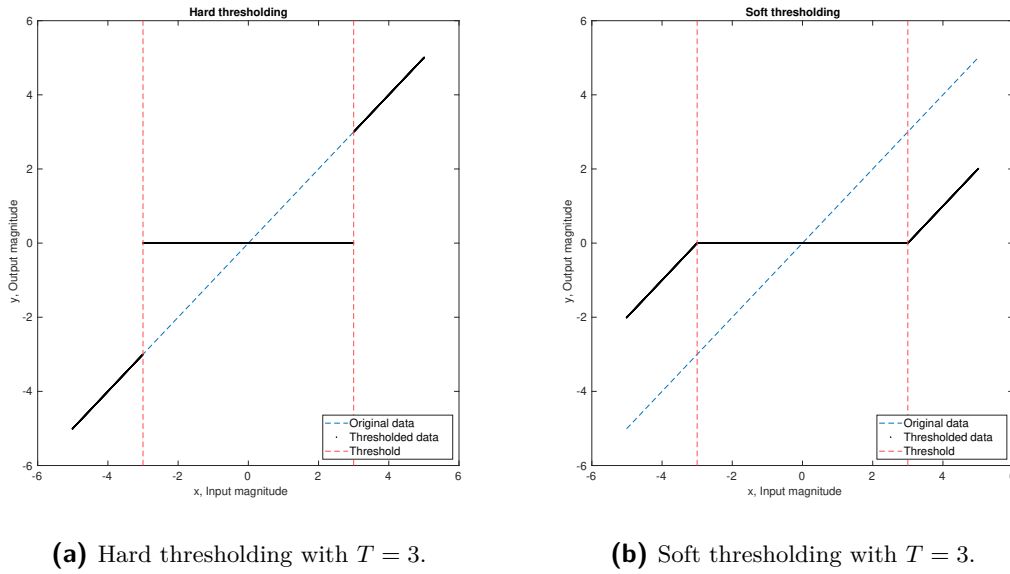


Figure 2-13: Wavelet thresholding methods.

2-4-2. Filtering

2-4-2-1. Adaptive thresholding

For Fourier transform based methods, periodic noise components are often identifiable visually in the form of localized peaks, as explained in Section 2-2-3. A simple filter would be to define a scalar magnitude threshold in the frequency domain, above which coefficients are considered to be noise and should be erased. In practice, this may work in simple cases, but often a more advanced thresholding algorithm is necessary to adaptively set a threshold and consider local or global information on the signal to be able to work in a more robust manner.

As an example Sur et al. [72], [74] used global information by segmenting the image to a number of sub-images and calculating the mean power spectrum of these patches. They then used the power-law distribution, which states that the power spectra of natural images are inversely proportional to the square of the frequency, also described as $1/f^2$ [20]. This means that when visualizing the power spectrum of a natural image on a logarithmic scale, it should follow a linear trend [20]. Sur et al. [72], [74] fitted a robust linear regression model to the power spectrum and defined a 3σ confidence interval around it to adaptively threshold components based on their frequency. Any frequency components over the 3σ line would be classified as noise.

A more local approach by Ionita et al. [82], [84] is to assume that the maximum of the power spectra are expected to be monotonically decreasing as the radial distance from the zero-frequency component increases. Their approach is to group amplitudes based on their radial distance from the DC component and make a vector $H(n)$ with the maximal values at each distance, then take the monotonically decreasing, minimum envelope of that vector as $T(n)$, see Equation 2-19. Whenever an amplitude is over this threshold, it is replaced by the mean of the amplitudes around the same radial distance from the zero-frequency component.

$$T(i+1) = \begin{cases} H(i), & \text{if } H(i+1) > T(i) \\ H(i+1), & \text{if } H(i+1) \leq T(i) \end{cases} \quad (2-19)$$

It is important to mention that high-magnitude, non-noise related components are often present in low frequencies, mostly in close proximity of the DC component. Removing these components changes the reconstructed image drastically, as they affect the mean of the image, but not the noise of interest. For this reason, many solutions [62], [70], [82], [84] define a "protected area" in this region, which is then excluded from noise detection and removal.

2-4-2-2. Frequency domain filters

Frequency filters are, often a combination of two basic filters. Namely, a low-pass filter that leaves frequencies under a specific edge frequency untouched, while attenuating everything beyond that edge frequency to zero. The second filter is a high-pass filter that attenuates low-frequency components while retaining high-frequency ones. The frequency where the change in behaviour occurs is called the cutoff frequency.

Frequency domain filters, such as the band-pass and band-reject filters are created by combining these two basic filter types. For example, a band-pass filter is created when a low-pass and a high-pass filter are connected in series, with the high-pass filter having a lower cutoff-frequency compared to the low-pass filter. A band-reject filter, also called band-stop filter is created when the band-pass filter is inverted, or by connecting a low-pass and high-pass filter in parallel, here the cutoff-frequency of the high-pass filter is higher than that of the low-pass filter. When looking at the frequency domain representation of these filters in 2D, their frequency response looks like o-shaped disks representing pass or reject areas.

Notch filters are a combination of the aforementioned filters where signals are passed or rejected in a specific neighborhood around a central frequency, individually defined for each dimension. This behavior results in pass-through or rejection patches, rather than disks in the frequency domain. This filter provides excellent control over which frequency components to keep and reject and is commonly used for post-processing images corrupted by periodic signals. More detailed information on these filters can be found in Chapter 4 and 5 of [81].

An example of each of the described filtering methods can be seen in Figure 2-14. In Figure 2-14c the effect of the DC component is visible, also discussed in Section 2-4-2-1. Comparing Figure 2-14b and 2-14c, it is also apparent that edges and fine details are kept with high-frequency components, while base colors and approximate shapes are preserved with the low-frequency components. It is interesting to note that the low-pass filtered image in Figure 2-14b only contains 0.03% of the original data, while Figure 2-14c contains 98.6% of the original data in the frequency domain.

Figure 2-14d rejects a band around where the noise component is, but also filters certain details of the image, while Figure 2-14e shows the exact band rejected. Figure 2-14f shows the notch filter that has been tuned almost to the exact frequency of the noise, and apart from the ringing phenomenon along the borders, it removes most of the noise while retaining the most fundamental image details.



(a) Original image, corrupted with periodic noise.



(b) Low-pass filtered image.



(c) High-pass filtered image.



(d) Band-reject filtered image or high-pass filtered image, with protected DC component.



(e) Band-pass filtered image.



(f) Notch-filtered image.

Figure 2-14: The effect of different basic filtering methods in the Fourier-domain.

Frequency-domain filtering is a simple, yet very effective method for the removal of noise components, that separate out well in the frequency domain. However, it can also remove genuine signal, that resides in the area of effect of the filters [60]. In more complicated cases, the removal of such signal components can partially corrupt the final result significantly, as much of the useful information may be concentrated in a handful of components. For this reason, more sophisticated techniques have been proposed for image restoration. For example, Sur et al. [72], [74] used a series of notch filters after classifying Fourier-coefficients as either noise or true signal. Even though this aggressively removed some information concealed by the noise, the results show that as long as the classification is accurate, the information-loss is minimal and good results can still be achieved by the denoising algorithm.

2-4-2-3. Noise component attenuation

Another option for restoration is that instead of completely removing (attenuating to zero) components that are identified as corrupted by noise together with potentially a portion of genuine signal, component are replaced by estimates of the expected value of the component. This approach attempts to remove the noise component, while retaining or at least approximating the noise-free value of the genuine signal.

This can be relatively easy and accurate in the case where the imaging can be done on a uniform, flat characterization sample, which provides a good baseline for the method noise of the system, as mentioned towards the end of Section 2-2-3. Unfortunately, such characterization runs are often not available or practical, as many datasets consist of images made on a multitude of differently calibrated machines and method-induced noise

can change and evolve over time, as well as a result of environmental effects, such as temperature and humidity.

Smart estimations can still be made, even though the uncertainty around the results tends to grow. One notable application of such an approach can be found in Sidorov et al. [32], where they suppress Moiré patterns appearing due to film-to-video conversion and apply it to real video sequences. In their algorithm, they take the Fourier transform of the image, then create a protective area around the DC value and its neighboring frequencies. Next, they choose a tuning parameter and a threshold, which is the midpoint of the maximal and minimal magnitude, above which they replace the original magnitude with the median of the non-protected values and then recombine the altered magnitude spectrum with the original phase and invert the Fourier transform.

Even though this algorithm removed a significant portion of the noise, there are visible residuals and filter-introduced signals even after careful tuning of the threshold. These residuals are likely present due to the aggressive truncation to the median that may overshoot the actual amount of attenuation required for a large number of coefficients closer to the protected area. Another shortcoming of this approach is the manual tuning, which makes it difficult and impractical to use it in an automated pipeline. Apart from the tuning parameter, the threshold is dependent on the outliers, running the algorithm multiple times in sequence with the same parameters yields different results.

The solution applied by Ionita et al. [70], [82] is to replace the identified noise components by the mean of the non-zero, non-noisy elements at the same radial distance to the zero-frequency component. This way the Fourier coefficients are not entirely rejected, as it would happen with a notch filter, but the noiseless value is estimated based on the information extracted from the image itself.

2-5. Image restoration

The primary goal of image restoration is to recover the true image from one that has been corrupted in a certain way. Most often, this is a highly objective process, as long as the type of corruption and the goal of the restoration are defined adequately. It is often achieved by defining, modeling, and then inverting the degradation process [81]. In this Chapter, we explore an approach that can mitigate the effects of artificial patterns introduced by filtering and look at a range of image quality assessment metrics. Furthermore, a comparison of strengths and weaknesses of these metrics is provided for specific applications found in the literature.

2-5-1. Total variation minimization

In the case of frequency domain filtering, it is possible for essential image features, such as line segments and edges, to be parallel to the direction and frequency of the periodic corruption that needs to be removed. An illustrative example would be noise that aligns with a mast of a ship or in the medical case, imagine a vertical leg CT, where the directionality of noise aligns with the direction of a bone.

As explored by Sur [73], spectrum interpolation using total variation minimization aids the mitigation of such artificially introduced ringing artifacts. It is important to note that even though the results are promising, the effectiveness of the method heavily depends on the regularization parameter, which has to be tuned for the specific application [22]. If this parameter is too high, the interpolation is ineffective. If it is too low, unwanted smoothing of the spectra occurs, which may hide important features, in addition to the

artifact at hand, which can be detrimental in the case of medical images, as explained in Section 2-2.

2-5-2. Metrics

Comparison of different algorithms and their optimal parameter settings is far from trivial. For example, in denoising, often one characteristic improves while another deteriorates. Determining if one algorithm or parameter setting is better than another is heavily dependent on the optimization goal. On top of that, quantifying these differences is far from trivial and adds another layer of complexity to the problem.

With a small number of signals or images to compare, assessing quality of image restoration and quality is often done manually by a domain expert. When the number of observations and thus comparisons grow, a consistent, automatic quality assessment method for image restoration is required, providing a mathematical definition of "distance" that lines up with what the expert would use implicitly. This is the goal of defining a metric. For the comparison of algorithms, it is vital to use a metric that is closely tied with the optimization goal. Otherwise, the results and perceived improvements may not be substantial. A range of image quality assessment metrics are introduced and a comparison of the strengths and weaknesses are provided.

2-5-2-1. Signal-to-noise ratio

In cases where the true, clean signal or image - denoted $s[n]$, where n is the number of elements in the signal - is available, in addition to the corrupted signal - denoted $x[n]$ -, the signal-to-noise ratio (SNR) is one of the metrics that can be used to measure the effectiveness of denoising. This metric is the ratio between the signal power and the noise power, as seen in Equation 2-20. It is commonly represented on the decibel (dB) scale. Measuring the effectiveness of denoising can be done, by comparing the SNR score of the unaltered, corrupt signal, $x_1[n]$, and the denoised signal, $x_2[n]$, with the same clean signal, $s[n]$. It can be used to compare the effectiveness of different denoising algorithms in the same manner. A higher SNR score is better, as it indicates a closer resemblance to the clean signal.

$$SNR = \frac{\sum_{i=1}^n s[i]^2}{\sum_{i=1}^n (x[i] - s[i])^2} = \frac{P_{signal}}{P_{noise}} \quad (2-20)$$

2-5-2-2. Peak signal-to-noise ratio

Peak signal-to-noise ratio (PSNR) is a modification of the SNR. It is the ratio of the peak intensity of the signal and the mean squared error, as shown in Equation 2-21. Its use cases are more content-specific. It is more sensitive to the bandwidth and dynamic range of the signal. Furthermore, it focuses more on the high-intensity performance of an algorithm. These qualities make this metric great for evaluating compression algorithms. It is also commonly applied to image denoising methods the same way as discussed in Section 2-5-2-1.

$$PSNR = \frac{\max(s[n])^2}{\frac{\sum_{i=1}^n (x[i] - s[i])^2}{n}} \quad (2-21)$$

2-5-2-3. Structural similarity index

According to Wang et al. [38], for applications where images are to be viewed or used by humans, the only correct way of quantifying quality and denoising performance is by subjective evaluation by an expert. In reality, this is expensive and time-consuming but still leads to better results than using traditional mathematical error metrics, such as mean squared error (MSE), SNR, or PSNR, which are often not suitable for evaluating tasks revolving around human perception, as in telecommunication and in medical imaging.

The structural similarity index (SSIM) is considerably different from the aforementioned methods, as it does not measure absolute differences but compares the signals using variety of properties derived from the human visual system [38]. This results in a perception-based, comparison of structural information. The SSIM algorithm makes comparisons on three different levels, namely luminance, contrast, and structure. All three of these comparisons are part of the final similarity measure. However, they are relatively independent, as changes in luminance, for example, do not propagate to a substantial difference in contrast and structure. The exact definition of the metric is not reproduced here, but can be found in [38].

2-5-2-4. Blind/Referenceless Image Spatial Quality Evaluator

In many practical cases, clean, ground truth images without corruption to compare against are not available, either because corruption happens during acquisition (e.g. in medical imaging techniques explored in Section 2-2), or because acquiring a clean image would require expensive hardware and can be a cost and time-intensive process. In these cases SNR, PSNR, and SSIM, all of which require a reference to compare against, are not usable for comparing and ranking different denoising algorithms or different parameter settings. Instead, Anish Mittal et al. [61] proposed a denoising performance metric, named the Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE), which functions without the need for a clean reference image. BRISQUE is a no-reference image quality assessment metric that does not use corruption-specific features, e.g. ringing or blur, but instead uses natural scene statistics (NSS). Specifically, it uses locally normalized luminance coefficients in order to quantify the deviation from "naturalness". BRISQUE assumes that the undistorted image is a natural image, where the probability density function of the normalized pixel intensities is close to a Gaussian distribution. They have shown that BRISQUE is competitive with respect to present-day non-corruption-specific no-reference Image quality assessment algorithms. The exact definition of the metric is not reproduced here, but can be found in [61].

Even though this approach does not focus on quantifying the amount of quasiperiodic patterns or lack thereof, it may be sufficient to quantify which denoising algorithms produce results that are closer to a natural image. It is important to mention that biomedical images are often not natural images, as described in [20]. The expectation is, that a noise-free image adheres to the NSS more than one corrupted by quasiperiodic patterns. If that is the case, this would be sufficient for the relative comparison of algorithms.

2-6. State-of-the-art

This section provides a brief summary of the state-of-the-art algorithms and methodologies available for characterizing quasiperiodic noise components and their removal, we also discuss any assumptions made about the corruption and the underlying noise-free signal, as well as, highlights of their most promising feats and their potential shortcomings.

2-6-1. ARPENOS

An algorithm created by Frédéric Sur and Michel Grédaic [73], [74], ARPENOS, is a quasiperiodic denoising algorithm that provides an automated approach to notch filtering and noise component labeling based on the expected power spectrum of a natural image.

Their most important assumption comes from the power law of natural images which states, in simplistic terms, that the logarithm of the power spectrum as a function of the logarithm of the radial distance from the zero-frequency component is approximately linear. Although this assumption seems to break down for low and high-frequency components, it remains robustly applicable for the majority of the power spectra. Furthermore, they assume that quasiperiodic corruptions are the only consistent periodic pattern throughout the image. They segment the image into a large number of overlapping patches, or "imassettes", and take the harmonic mean of the power spectra. Then, the peaks of the resulting aggregate power spectra are good descriptors of the quasiperiodic corruption. The size of the patches used determines the effective frequency bandwidth of the resulting power spectrum, and thus the upper bound of the detection horizon of the algorithm. The main steps of the algorithm are shown on Figure 2-15. The aggregate power spectrum is then transformed to a one-dimensional description, as a function of the radial distance from the central component. Robust linear regression is applied to the power spectrum on a log-log scale. Since the linearity assumption is seriously violated for the lower frequency range, points in the lower frequency range are excluded from the regression procedure, commonly removing the lowest 10-20% of the power coefficients. Components are labeled as outliers when their power is higher than the regression line plus three standard deviations of the same regression line. This is effectively an outlier labeling approach based on the z-score. The detected outliers are then projected back to their original locations in the two-dimensional coordinate system. Based on the outlier map, notch filters are automatically designed to remove the outliers. Finally, the filtered power spectrum is combined with the original phase and the inverse Fourier transform is applied to obtain a denoised image in the original measurement space.

Manually tunable parameters include the size of patches, the step size and overlap of the patches, the lower and upper frequency bounds of the regression process, the lower frequency bound for outlier detection, and the number of standard deviation multiples that characterize an outlier. Furthermore, smoothing through total variation minimization is also possible, which helps to reduce patterns caused by Gibbs phenomena. However, this adds yet more parameters, such as the number of maximal iterations and epsilon, the regularization parameter.

The ARPENOS algorithm is relatively simple, its computational costs are relatively low, and its denoising performance is high. However, the assumptions made about the underlying image limit its effectiveness significantly, especially in the low-frequency range. Furthermore, a number of parameters have to be tuned individually for good results. These parameters are highly image-specific, and have to be re-tuned for different images.

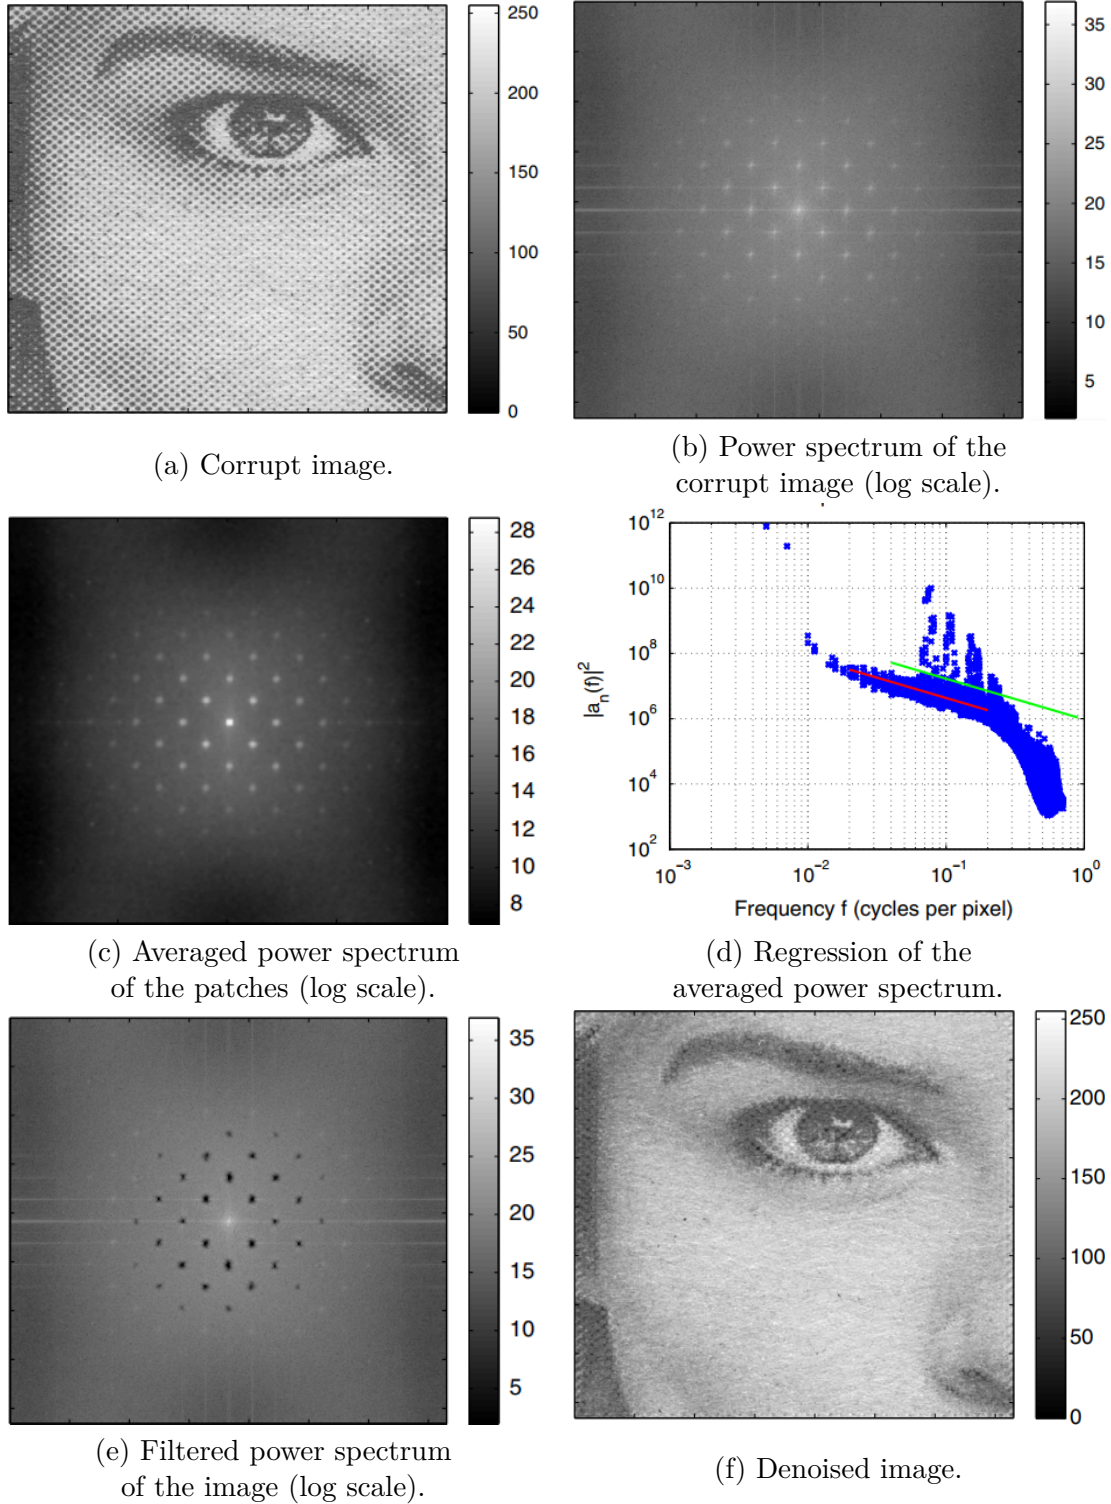


Figure 2-15: Main steps of the ARPENOS algorithm [74]. In the one-dimensional image, the red line indicates the regression line, its horizontal limits indicate the borders of the linear region. The green line indicates the upper bound that is three standard deviations over the trendline, horizontally, it shows the lower-frequency detection limit to be at 0.04 cycles/pixel. Source: [74], Figures 16,17, rights to reuse requested from the Society of PhotoOptical Instrumentation Engineers (SPIE).

2-6-2. ACARPENOS

An algorithm created by Frédéric Sur [72], ACARPENOS, is a quasiperiodic denoising algorithm that provides an automated approach to notch filtering and noise component labeling based on the idea that features of interest are not likely to be caused by a random background process and that they can be identified based on the number of false alarms (NFA).

This algorithm is based on ARPENOS and shares most of its steps. However, its outlier detection is altered, compared to the previous method. The main steps are illustrated in Figure 2-16. The labeling approach is not based on the expected power spectrum of a natural image, but instead on a probabilistic approach to independent and identically distributed random variables. Outliers are labeled based on the probability of occurrence in a local population. The image is segmented into a large number of overlapping patches, or "imageettes", similar as before, where P denotes the number of patches. Then the minimum of individual coefficients are taken from the population of P values. This serves as the estimate of the background process. Subsequently, the power spectrum is radially binned using a number of equidistant, concentric disks, an approach based on the power law of natural images [20]. This power law states that the expected value of power coefficients of natural images are dependent on frequency, such that on a log-log scale this dependence is linear. In the ACARPENOS case, linearity is not assumed, but radial dependence on frequency is. The distribution is empirically estimated for each ring, which serves as the estimate of the noise-free distribution. Within a given ring, only a small number of coefficients may belong to spikes caused by quasiperiodic noise. Then, the expected number of high values or NFAs is calculated for each coefficient, based on the background process and the estimated distribution. If the NFA for a certain coefficient is higher than the set threshold, it is labeled as an outlier. The detected outliers are then extrapolated to their original locations in the two-dimensional coordinate system. Based on the outlier map, notch filters are automatically designed for the outliers. Finally, the filtered power spectrum is combined with the original phase and the Fourier transform is inverted. This results in a denoised image.

Manually tunable parameters of ACARPENOS include the size of patches, the step size and overlap of the patches, the frequency limit that protects the zero-frequency and neighboring coefficients, and the allowed threshold of the NFA. Furthermore, smoothing through total variation minimization is also possible, which can help to reduce patterns caused by Gibbs phenomena. However, this again adds more parameters, such as the number of maximal iterations and epsilon, the regularization parameter.

The ACARPENOS algorithm is an interesting probabilistic approach, with increased computational costs compared to ARPENOS, but these are still reasonable, and denoising performance is relatively high. Furthermore, a number of parameters still have to be tuned individually for good results. Since these parameters are highly distribution-specific, they usually have to be tuned for different images regularly.

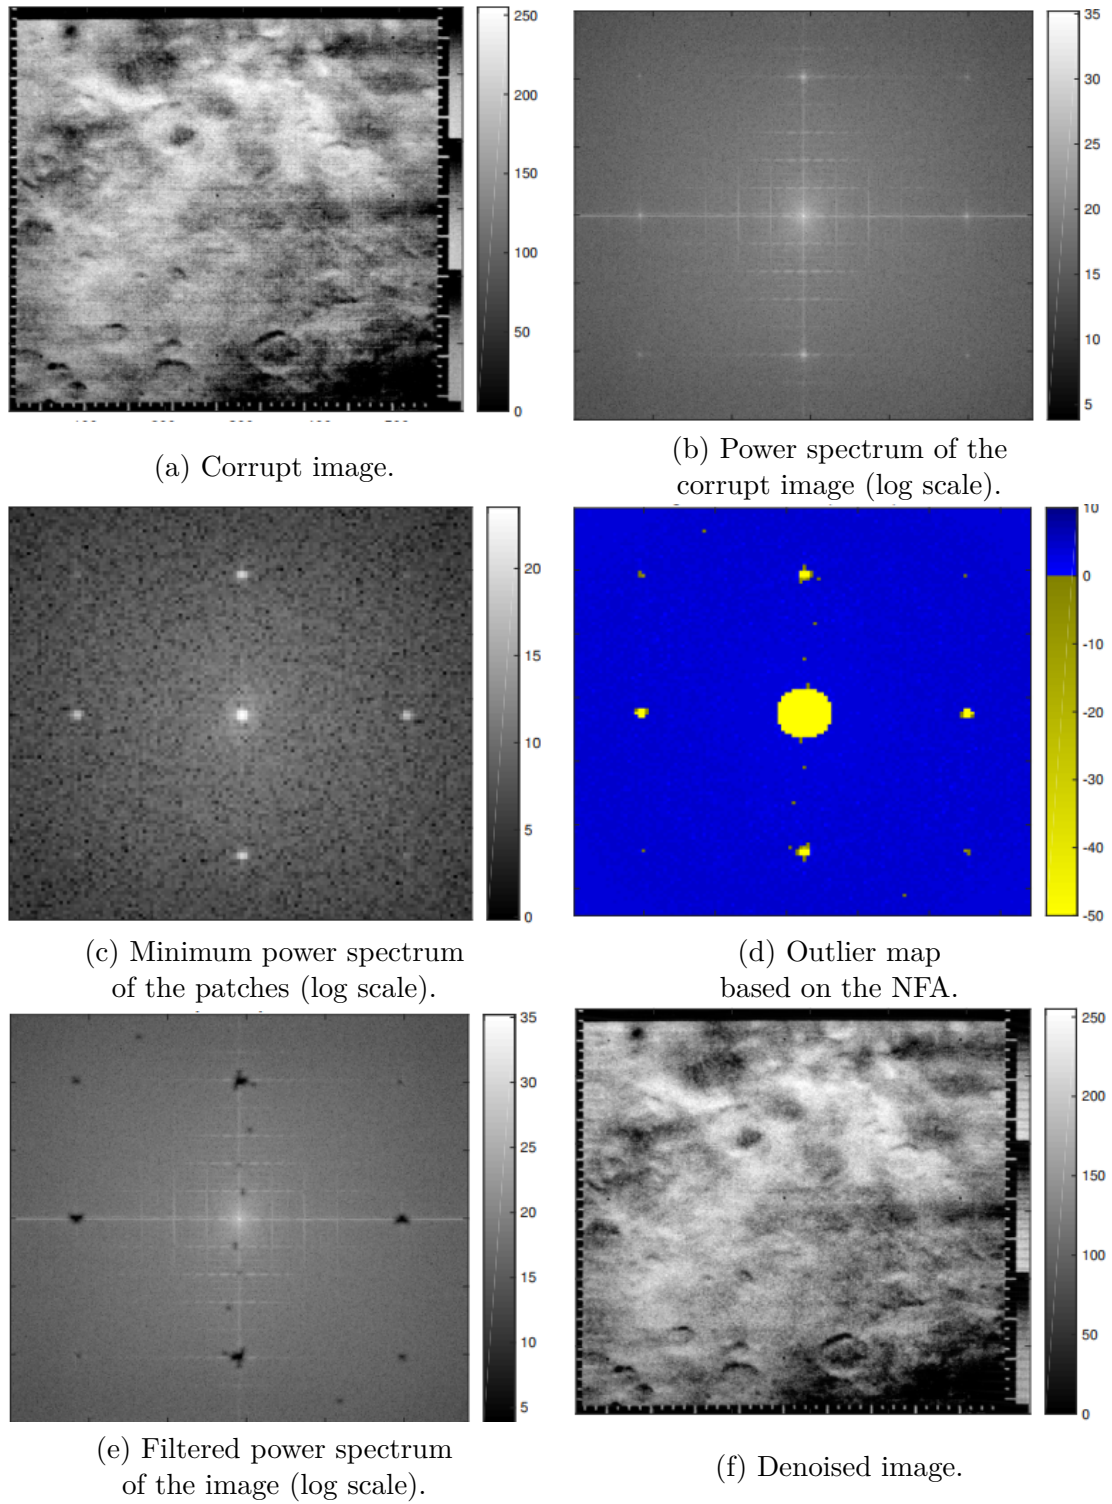


Figure 2-16: Main steps of the ACARPENOS algorithm. In the outlier map yellow indicates outliers, which are an unexpectedly high NFAs. The central area is protected from filtering. Source: [72], Figure 1, ©2015 IEEE.

2-6-3. IONITA


An algorithm created by Marius G. Ionita and Henri George Coanda [70], is a quasiperiodic denoising algorithm that uses radial patterns in the power spectrum to label outliers, and then attenuates them to the average of the components with the same radial distance.

Their most important assumption is that the maximal values of the power spectra are expected to be monotonically decreasing as the radial distance from the zero-frequency component increases. The power spectrum is transformed to one dimension, as a function of the radial distance from the central component. From this one-dimensional representation of the power spectrum, a vector $H(n)$ is created with the maximal values at each distance, where n is the radial distance to the DC component. Then the monotonically decreasing, minimum envelope of vector $H(n)$ is taken according to 2-22 as $T(n)$. Each amplitude is compared to $T(n)$ at the respective radial distance. If the difference is larger than a predefined threshold, the component is replaced by the mean of the components at the same radial distance. A protection area is created around the zero-frequency component, components within this region are excluded from attenuation. The main steps of the algorithm are shown in Figure 2-17.


$$T(i+1) = \begin{cases} H(i), & \text{if } H(i+1) > T(i) \\ H(i+1), & \text{if } H(i+1) \leq T(i) \end{cases} \quad (2-22)$$

Manually tunable parameters of IONITA include the radial frequency limit that protects the zero-frequency and neighboring coefficients, and the threshold. Recommendations for choosing the threshold and the exact thresholds for the example image were not provided in the paper.

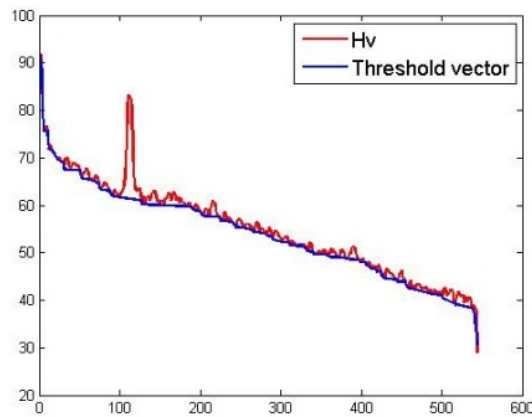

The IONITA algorithm is a straightforward approach, that seems to achieve a relatively high denoising performance with a relatively low number of tunable parameters.




(a) Corrupt image.



(b) Power spectrum of the corrupt image.

(c) The radial maximum vector, $H(n)$ and its minimum envelope (log scale).

(d) Filtered power spectrum of the image.



(e) Denoised image.

Figure 2-17: Main steps of the IONITA algorithm [70]. The specific threshold used, the outlier map, and the scales of the power spectrum are not shown as they were not provided in the original paper. Source: [70], Figure 2,4,6, ©2015 IEEE.

2-7. Summary

In this chapter, a comprehensive comparison between decomposition and domain transformation methods is given in Section 2-3. The most prominent filtering and identification methods are described in Section 2-4. Finally, a comparison of the available metrics listed in Section 2-5 is carried out, focusing on periodic noise removal of biomedical images. With regards to domain transformations, staying in the spatial-domain is often insuffi-

cient in the majority of the cases, as Sidorov et al. [32] explored. Parametric modeling may work in simple cases, but tends to perform sub-optimally compared to the Fourier domain methods available. Fourier transform based methods are particularly suited for periodic and quasiperiodic noise removal [62], [70], [72], [74], [82]. Most of these use adaptive thresholding algorithms for outlier detection, based on local or global information, and then apply either notch filtering or noise component attenuation. While these algorithms seem particularly promising, all of them assume that the noisy components are reasonably well-localized, which is true for pure periodic noise, but often not the case for its quasiperiodic counterpart. Another concern is that most of these algorithms use SNR or PSNR as an evaluation metric, which may not be well suited for performance evaluation on biomedical images [45]. Wavelet-transform based approaches may handle a wide range of noise types and have been successfully applied to quasiperiodic noise by Ionita et al. [84]. Because of the more temporally aware decomposition, it handles quasiperiodic noise well. However, accurately identifying noise components in the wavelet subbands is more complex as they are not as well localized as in the case of a Fourier domain [55].

Statistical outlier detection methods are widely used. However, standard deviation based approaches should only be used when the underlying distribution is reasonably close to a Gaussian normal distribution and it should be avoided for e.g. Gaussian mixture distributions. Depending on the severity of the corruption, methods that are robust to outliers should be preferred, such as modified z-score [14] or interquartile-range based methods [42].

Wavelet-based soft and hard-thresholding methods are simple yet effective. However, they require manual tuning for good performance. Adaptive thresholding algorithms work well in both the Fourier and wavelet domain, particularly ones that use local information [62], [84]. These usually require tuning parameters that are less sensitive, such as the kernel-size in median filtering.

After proper identification and localization of noise components, basic filtering techniques, such as band-pass and band-reject, are often too aggressive and filter out a significant portion of the true signal. Successful applications of automatic notch filter design have been shown by Sur [74]. Notch filtering combined with spectral total variation minimization [73] was effective at mitigating the ringing effects caused by notch filtering, by smoothing out the frequency regions affected by the filter. The performance of this combination was only rivaled by noise component attenuation, where local or global information was used to estimate the noise-free coefficients [62], [82].

Most quality assessment metrics require a clean, ground truth signal, which is only available in case of synthetic data. Some metrics, like SNR and PSNR, focus on signal power while others, such as SSIM, focus on human perception. No metric was found that corresponds to the medical goals of denoising, without the involvement of medical experts, although some of these considerations can be found in Pižurica et al. [45]. On the other hand, BRISQUE, a no-reference score proposed by Mittal et al. [61] demonstrates an algorithm that does not require a clean reference image but instead uses scene statistics, in order to determine the "naturalness" of images. This may be used to rank different denoising algorithms by their effectiveness.

Some of the state-of-the-art algorithms are briefly introduced, namely ARPENOS [74], which is based on the power law of natural images and provides an automatic notch filter design in the Fourier domain. ACARPENOS [72], which is a probability based approach operating in the Fourier domain. It also provides an automatic notch filter design for the removal of detected outliers. Finally, IONITA [70] expects the maximal values of the power spectrum to be monotonically decreasing as a function of radial distance to the central

component. Components that significantly deviate from this radial pattern are marked as outliers and replaced by the mean of the components at the same radial distance.

To conclude, there is a wide variety of denoising algorithms, some for medical imaging, some for periodic noise cancellation, but the intersection of these fields is relatively scarce. The ideas, however, are very promising, domain transformation techniques such as the Fourier and wavelet transform are the basis of most effective algorithms. Matrix decomposition methods such as NMF and RPCA can help reveal hidden structure in the data and make separation and labeling of noise and signal of interest easier using adaptive thresholding methods. In terms of filtering, noise component attenuation methods seem to be most effective. However, if the estimate of the true value is not available, then notch filtering together with spectral smoothing using total variation minimization is also a suitable solution. The metrics used should be a combination of widely used metrics that are sensitive to different characteristics, so that their combination can mitigate their individual shortcomings and reveal interesting patterns in the algorithms' behaviour.

All state-of-the-art algorithms have shortcomings such as manual tuning of multiple parameters, insensitivity to low-frequency corruptions, and rather strong assumptions about the underlying noise-free signal. We did not find an algorithm in the literature, that is entirely data-driven and is able to reliably identify, quantify, and eliminate quasiperiodic patterns, specifically focusing on biomedical images. The development of a new algorithmic framework that satisfies these goals is therefore necessary.

Chapter 3

Methods

In this Chapter, a new framework is introduced, named Quasiperiodic Image Denoising (QID). Within this framework, two novel denoising algorithms are implemented, both operating in the Fourier domain. One algorithm is based on robust principal component analysis (QID-RPCA) and the other uses the normalized median of absolute differences (QID-MADN).

This section provides more detail on the specific steps taken in the novel algorithms, QID-RPCA and QID-MADN, and elaborates on their usefulness towards an unsupervised, fully data-driven denoising framework. The steps can be categorized into three groups: preprocessing, quasiperiodic noise component labeling and filtering, and image reconstruction.

Figure 3-1 presents the main steps of the framework. Figure 3-1a shows the image, that has a visible high-frequency corruption. Figures 3-1b and 3-1c show examples of the preprocessing steps. These preprocessing steps are e.g. transformation to the Fourier domain, specifically the power spectrum of the image and radial binning through histogram equalization. Figures 3-1d and 3-1e show outlier detection and attenuation of these outliers. Figure 3-1f shows the reconstructed, denoised image.

First, an overview of the pipeline is provided, where each step is introduced briefly. Then, a flowchart is presented, which is used as a reference throughout this chapter. Each step is explained in detail. Adjacent blocks of the flowchart are shown at the beginning of each step.

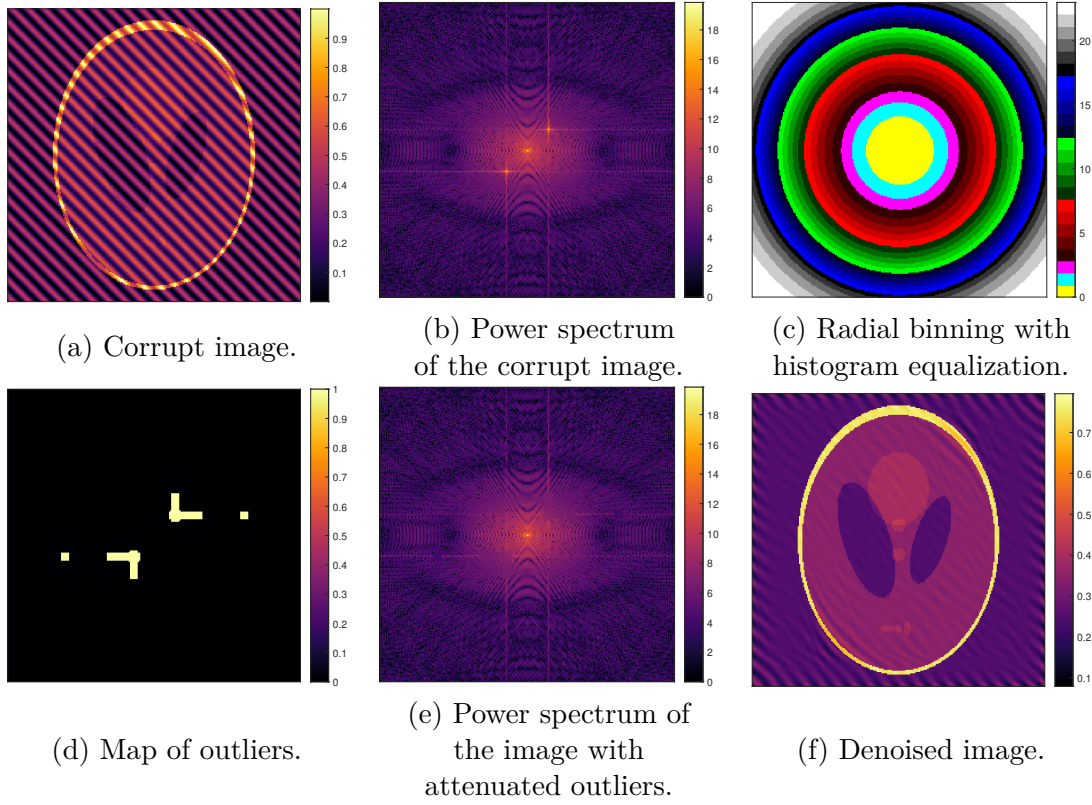


Figure 3-1: Examples of the most important steps of the QID denoising pipeline.

3-1. Pipeline

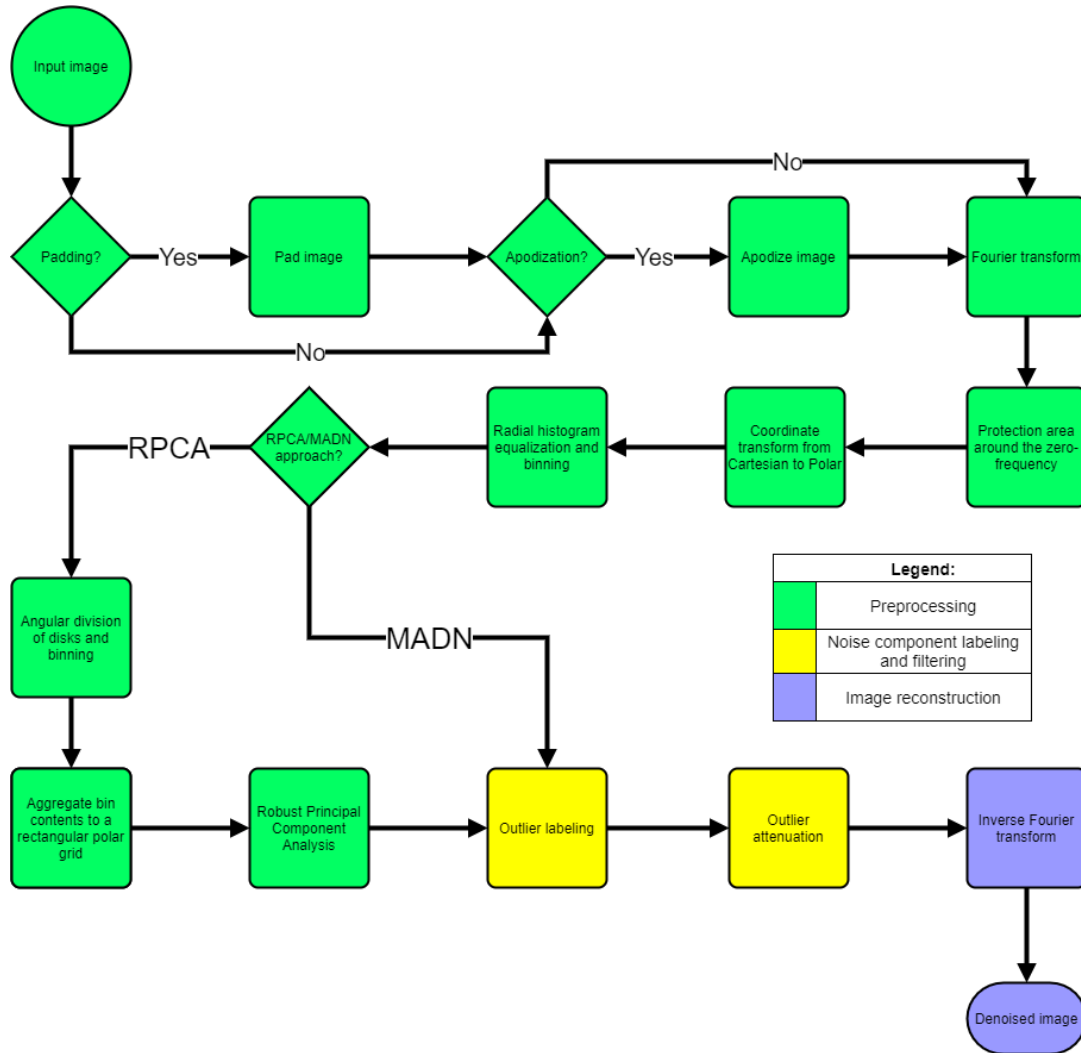


Figure 3-2: Flowchart detailing the steps of each version of the algorithm.

The steps of QID can be categorized into three groups: preprocessing, quasiperiodic noise component labeling and filtering, and image reconstruction. After a brief introduction of the steps, a block diagram representation is presented on Figure 3-2.

The steps are as follows:

- Preprocessing:
 1. Padding and apodization: padding artificially increases the number of pixels. This can be used to artificially create a finer description of the Fourier representation, while keeping the frequency bandwidth the same. Windowing can smooth the ends of a period such that no abrupt jumps occur. This also modifies the resulting Fourier representation and can alter the underlying biological information. This can be used to e.g. reduce the negative effects of the Gibbs-phenomena, if applied with consideration.
 2. Fourier transform: it approximates any image domain signal as a sum of simple, periodic functions. Periodic and quasiperiodic patterns can be characterized

well in the frequency domain, and thus this domain transformation step is a central part of this analysis.

3. Protection area for the zero frequency and adjacent components: changes to the zero frequency component and some of its neighboring coefficients affects the image mean significantly, while quasiperiodic patterns rarely affect this area. Altering them has detrimental effects to the reconstruction, thus a protective mask is built to exclude the zero and low-frequency components from filtering.
 4. Polar coordinate system transform: Fourier coefficients of biomedical images show a similarity in magnitude for coefficients at a similar radial distance from the central zero frequency component. Transformation to the polar coordinate system is needed to provide an expected value for the components identified as noise.
 5. Radial binning using histogram equalization: the Fourier coefficients are binned to a number of radial disks. Through histogram equalization, all bins have a similar number of coefficients contained. Histogram equalization is required to ensure that the eventual labeling is done on a sufficiently large, but equal population among bins, which is representative of the radial distance.
 6. Angular division of disks: in the case that QID-RPCA is used, the disks are further divided into angular bins. This is necessary for mapping polar values to a rectangular grid.
 7. Mapping polar values to a rectangular grid: in the case that QID-RPCA is used, the components in each angular and radial bin are aggregated. This forms a rectangular grid. This step ensures, that RPCA is able to reliably find a sufficient low-rank and sparse representation.
 8. Robust Principal Component Analysis: in the case that QID-RPCA is used, RPCA is performed on the rectangular grid of aggregated polar values. This separates the low-rank, or expected value for a certain bin, from the sparse, or corrupt value.
- Noise component labeling and filtering:
 1. Outlier labeling: in the case that QID-RPCA is used, labeling is done using z -scores on the sparse matrix. If the QID-MADN variant is chosen, labeling is done using the modified z -score. This step marks the coefficients that are most likely to come from a noise source and not from a biological origin.
 2. Outlier attenuation: in the case that QID-RPCA is used, corrupt components are attenuated to their corresponding value in the low-rank matrix. This is done for the Fourier components in the original coordinate system and not in the rectangular representation of polar values. If the QID-MADN variant is chosen, corrupt components are attenuated to the radial disk's median. Corrupt patterns are effectively filtered in this step and are replaced by the best estimate of their noise-free value.
 - Image reconstruction:
 1. The modified magnitude components are combined with the original phase. The inverse Fourier transform is applied, resulting in a reconstructed image in the original domain.

3-2. Apodization and padding




Figure 3-3: Pipeline blocks relevant to padding and apodization.

The 2-D discrete Fourier transform is focused on periodic signals. The window in which periodicity can be captured is the full image width and height. Whenever discontinuities occur between the horizontal or vertical pairs of edges of the image, the perfect Fourier representation tries to capture this perceived abrupt discontinuity. Practical applications, however, are only an approximation of this perfect periodic representation and have a limited resolution to represent such abrupt jumps. This often results in the introduction of artificial patterns, e.g. the Gibbs-phenomenon, near these discontinuities [64]. As a result, when images with such discontinuities are denoised using the limited-range Fourier transform, noise patterns may appear towards the edge of the image. An example of this behaviour is shown for a simple one-dimensional time signal on Figure 3-4.




Figure 3-4: An example of the Gibbs-phenomenon. It appears at discontinuities of the time or image domain when the discrete Fourier transform approximates the theoretical Fourier representation with a limited range.

In order to mitigate this effect, apodization or more commonly known as windowing, may be used. Windowing smooths the differences at the edges, avoiding the introduction of

a perceived discontinuity. However, it also alters the original signal [81]. Depending on the affected area, it can lead to significant true signal loss, if applied without careful consideration. In the case of biomedical image denoising, even minimal loss of true signal can have detrimental effects, e.g. misdiagnosis.

The area of effect for windowing has to be set such that no relevant information is contained near its edges. In some instances, this reduces the possible area of effect to a very narrow frame around the true signal. However, certain windowing functions, like the Tukey window, can restore periodicity in such a region [80]. The resulting change can still be abrupt and artificial patterns may appear, possibly elsewhere. For a more complete mitigation of the problem, a slow, gradual change is required. Image padding is a possible solution to this. Equation 3-1 describes the one-dimensional Tukey window, also called the tapered cosine function. The edges of this window function are phase shifted halves of the cosine functions, while the central area is a rectangular window. $w(x)$ is the window function, where x is an N -point linearly spaced vector, r is a tunable parameter that represents the proportion of the cosine function within the whole window. If $r = 0$, the window is equivalent to the rectangular window. If $r = 1$, it is equal to the Hann window. An example of denoising with an aggressive Tukey window, with $r = 0.08$, can be seen on Figures 3-5a and 3-5b. Compared to Figures 3-5g and 3-5h, where no windowing is applied, the patterns around the edge of the image are attenuated. However, new, artificial patterns are introduced as well.

$$w(x) = \begin{cases} \frac{1}{2}(1 + \cos \frac{2\pi}{r}(x - \frac{r}{2})), & 0 \leq x < \frac{r}{2} \\ 1, & \frac{r}{2} \leq x < 1 - \frac{r}{2} \\ \frac{1}{2}(1 + \cos \frac{2\pi}{r}(x - 1 + \frac{r}{2})), & 1 - \frac{r}{2} \leq x < 1 \end{cases} \quad (3-1)$$

Through image padding, the sides of the image are expanded with artificial values to increase the area where no true signal is contained. A gradual transition can be made through apodization while circumventing signal loss. Common choices for padding values are zeros, the image mean, or the mean of a particular border around the edge of the image. For medical imaging techniques such as in vitro microscopy, where the sample is placed on a plate before imaging, using the border for padding can result in a transition without any visible discontinuities. An added benefit of padding, when using the Fourier transform, is the appearance of more frequency bins, which results in a smoother, more detailed spectral representation. Even though the effective resolution of the spectrum does not change, padding comes at a price, the image increases in pixel size. Depending on the chosen padding size, computation costs can increase significantly, which is why careful consideration of such parameters is required. An example of denoising of a mean padded image with 50% of the width/height on each side is presented on Figures 3-5c and 3-5d. In this specific example, since the corruption pattern has a high amplitude, and it also affects the edges of the image, padding also introduces an abrupt discontinuity. Compared to Figures 3-5g and 3-5h, where no padding is applied, artificial patterns are introduced due to the abrupt change where padding starts.

Figures 3-5e and 3-5f show an example of the combined usage of padding, then Tukey windowing. The unaltered, rectangular portion of the window starts at the edges of the raw, non-padded image. In this case the artifacts introduced by padding are still visible, but the windowing did not introduce artifacts into the biological signal.

One of the main goals of the QID algorithms is to minimize the introduction of artificial patterns, while removing noise. Although Figure 3-5 portrays a heavily simplified setting, it aims to illustrate that introduction of artificial patterns through padding and window-

ing is relatively easy. Though certainly possible, setting such padding and windowing parameters in a non-intrusive, data-driven fashion lies outside the scope of this document. The exact parameter settings of both apodization and image padding can be highly specific to the image, the scene imaged, and the application at hand. For this reason, our default setting for QID is not to pad, nor to apodize the images. Our implementation contains padding and apodization as optional tools, but since data-driven operation is paramount for this thesis and these options require manual parameter setting to deal with scene-specific avoidance of introducing non-biological signals, the case studies shown in this thesis do not use them.

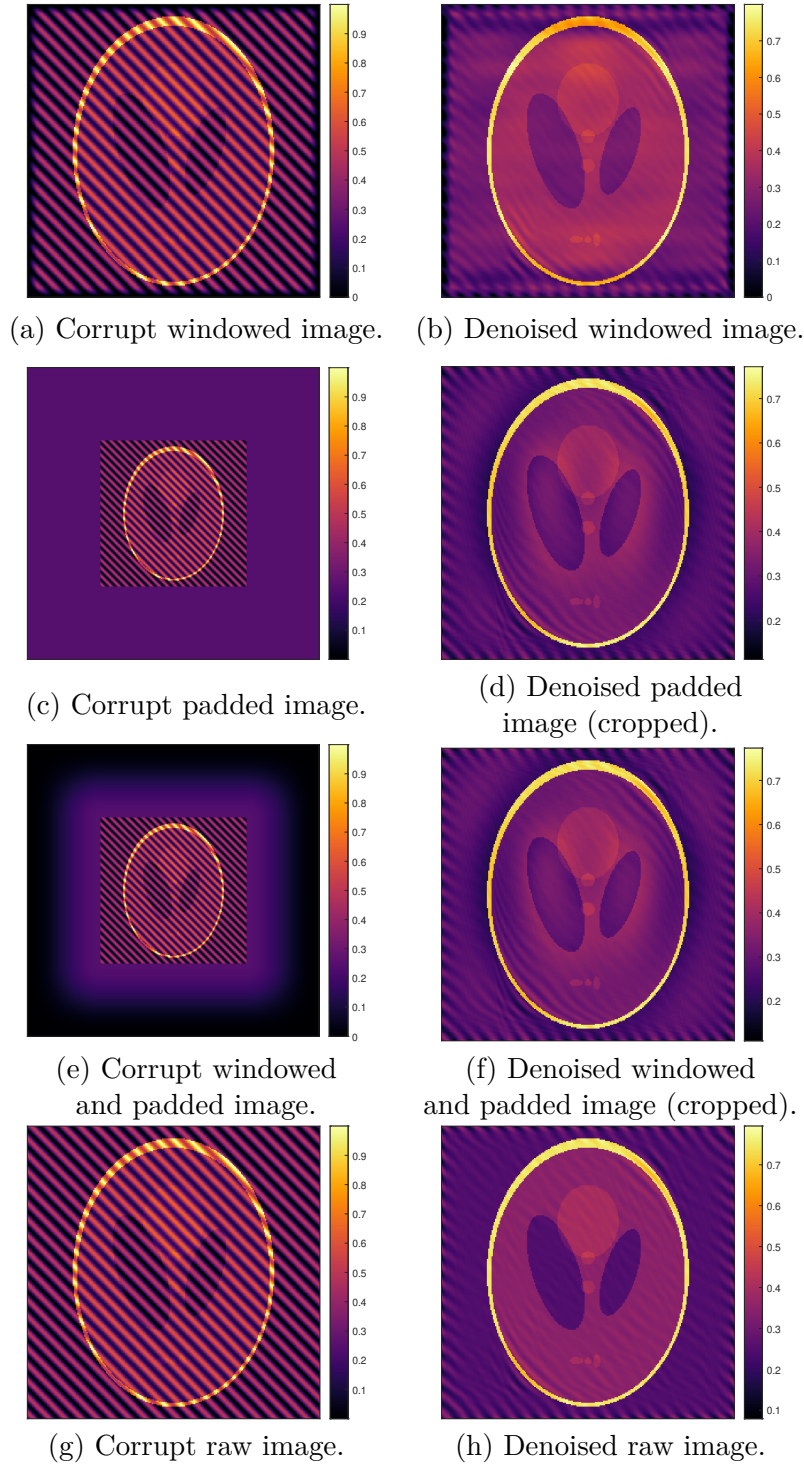


Figure 3-5: Examples of the effects of padding and windowing. Highlighting the emergence of artificial patterns.

3-3. Fourier transform

Details on the Fourier transform and analysis in the Fourier domain can be found in Section 2-3-1-2.

The 2-D discrete Fourier transform implementation we used is shown in Equation 3-2. For an $n \times m$ matrix in the spatial domain, X , and its discrete Fourier transform Y ,

$\omega_m = e^{-\frac{2\pi i}{m}}$ and $\omega_n = e^{-\frac{2\pi i}{n}}$ are complex roots of unity, where i is the imaginary unit, p and q are indices in the Fourier domain, j and k are indices in the spatial domain, p and j are indices that run from 0 to $m-1$, and q and k are indices that run from 0 to $n-1$.

$$Y_{p,r} = \sum_{j=0}^{m-1} \sum_{k=0}^{n-1} \omega_m^{jp} \omega_n^{kq} X_{j,k} \quad (3-2)$$

The resulting Fourier coefficients are complex numbers that can be represented by their magnitude and phase. In Equation 3-3, $a_{p,r}$ and $b_{p,r}i$ are real and imaginary numbers, respectively. $A_{p,r}$ is the magnitude and $\Phi_{p,r}$ is the phase.

$$Y_{p,r} = a_{p,r} + b_{p,r}i = \sqrt{a_{p,r}^2 + b_{p,r}^2} e^{i \arctan \frac{b_{p,r}}{a_{p,r}}} = A_{p,r} e^{i\Phi_{p,r}} \quad (3-3)$$

Further analysis will be done on the power spectrum of the Fourier coefficients, which is the square of the magnitude, as illustrated in Equation 3-4. This assists detection of outliers, as coefficients with an already high magnitude compared to its neighbors, are amplified even further. From this point on, coefficients of the power spectrum are used for the computations, even when referencing "Fourier coefficients". No changes to the phase are made. However, it is used in the final step when reconstructing the image with the filtered power spectrum coefficients and the original phase obtained in this step. Examples of the power and phase spectrum are shown on Figure 3-6.

$$P_{p,r} = a_{p,r}^2 + b_{p,r}^2 = A_{p,r}^2 \quad (3-4)$$

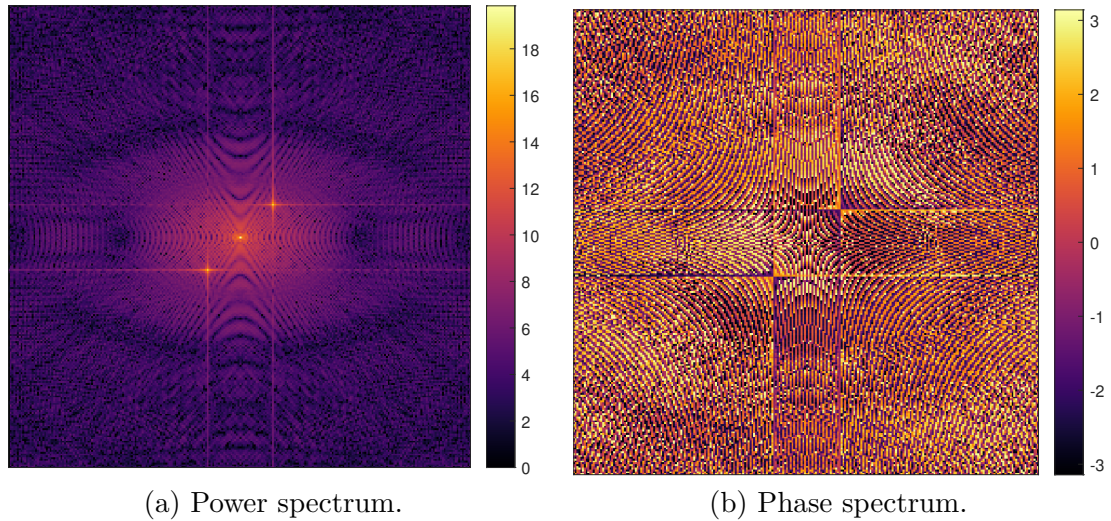


Figure 3-6: Examples of the Fourier transformed values of a corrupt image.

3-4. Protection area for the zero frequency and adjacent components




Figure 3-7: Pipeline blocks relevant to protection area masking.

The Direct-Current (DC) component represents the constant, time-invariant part of the signal or spatially-invariant part of the image. For one-dimensional signals, the mean amplitude of a signal can be described by the DC component, also commonly referred to as the DC value or DC offset. When transformed to the Fourier domain, the DC component is represented by the Fourier coefficient at 0 Hz for time-variant signals or 0 cycle/pixel for spatial signals such as images. In the case of the discrete Fourier transform, this is captured in the zero frequency bin, from this point on, we will use the "zero frequency bin" and "zero frequency component" interchangeably. The DC component is present in the time/image domain while the zero frequency component and zero frequency bin are its corresponding representation in the Fourier domain.

In the image domain, using a grayscale image as an example, the DC component represents the mean intensity of the image. In the 2-D Fourier transform of said image, the zero frequency component is contained in the first bin, or in the quadrant shifted case, the central bin of the Fourier domain representation. Quadrant shifting in the 2-D case separates the matrix into four quadrants at the central index of each frequency, creating half-spaces and swaps the half-spaces of the matrix along each dimension. This relocates the first element of the matrix to the center. Both of these are valid ways of representing the frequency content of an image and it is a matter of preference, throughout this document, the zero frequency component is contained in the central bin. Any change to the zero frequency component value directly affects the mean intensity of the image. The image mean is not altered by purely periodic corruption, which is why it is expected to remain unaltered during the denoising process. An example of what high and low-frequencies translate to in the image domain has been shown earlier, on Figures 2-14b and 2-14c. Quasiperiodic corruption and corruption with incomplete periods may alter the mean, but this effect is often negligible. However, depending on the resolution and accuracy of the Fourier representation, leakage to the neighboring coefficients may occur. Therefore, most state-of-the-art algorithms implement some form of protected area for the zero frequency component and its immediate surroundings. This protected area prevents alteration of these components, even when in unprotected circumstances these low-frequency coefficients would be labeled as outliers. Figure 3-8 shows the effect of changes made to the protection area. Figure 3-8a shows an adequately protected mean, while Figure 3-8b filters some of the components that affect the mean. This introduces a low-frequency pattern in all directions. Figure 3-8c shows the effects of removing the protection area completely. The reconstructed image without protection has approximately zero mean, while the ground truth image has not at all.

The ARPENOS, ACARPENOS, and IONITA algorithms handle the extent of this protected area using a user-set parameter. For ARPENOS, the radius of the protected area is usually set to cover a large area, as their linearity assumption based on the power-law distribution is violated for low-frequencies. Hence, setting the protected area radius to a

low value decreases ARPENOS' performance significantly. In Sur et al.'s experiments, the radius was often set to about 20% of the maximal possible radius [75]. Therefore, these algorithms may be unreliable or even altogether ineffective for low-frequency quasiperiodic corruptions.

As opposed to ARPENOS, ACARPENOS, and IONITA, the QID framework does not assume linearity or any specific trend in the Fourier coefficients. The protection radius can be set with considerably more flexibility. Empirical experiments suggest that a circular protection area with a radius of five to ten coefficients or values around the zero frequency component is sufficient for a variety of image sizes and corruptions. As a consequence, a radius of five coefficients is used for images where the shortest image dimension is under 512 pixels and its radius linearly grows until the shortest image dimension is 1024 pixels. In cases where leakage from the zero frequency component to the adjacent coefficients occurs, most of the power is still contained in these components. The benefits of this becomes apparent in Section 3-10, but, in short, as the number of pixels increase, the number of bins often increase as well. Automatically increasing the radius of the protection area for high-resolution images, with more and larger bins, retains the precision of the expected value for low-frequency bins. This simple modification increases denoising performance for high-resolution images.



(a) Denoising with the default radius of 5 frequency coefficients.



(b) Denoising with a radius of 2 frequency coefficients.



(c) Denoising with no protective area.

Figure 3-8: Examples of the effect of the protection area radius.

3-5. Cartesian to polar




Figure 3-9: Pipeline blocks relevant to polar coordinate system transformation.

The 2-D Fourier representation of an image in the Cartesian coordinate system decomposes the image into superpositions of both horizontal and vertical periodic functions. In the polar coordinate system, the horizontal and vertical components, x and y are translated to radial and angular components, r and θ , respectively. The radius describes the superimposed signal's frequency in terms of cycles/pixels. At the same time, the angle defines a rotation along the axes. Figure 3-10 shows an example of radial and angular coordinates of a power spectrum. The angular component is wrapped to $[0, \pi]$, to recognize the central symmetry of the two-sided Fourier transform.

$$r = \sqrt{x^2 + y^2} \quad (3-5)$$

$$\theta = \tan^{-1} \frac{y}{x} \quad (3-6)$$

Amplitude components of natural images with the same radius, but a different angular coordinate, are generally expected to be close in value, or at least in the same order of magnitude [20]. This assumption is broken by most periodic and quasiperiodic patterns. This phenomenon is exploited by many state-of-the-art algorithms to varying degrees [70], [72], [73], [82]. Analysis of components in the polar coordinate system provides a more intuitive description of what the expected value may be.

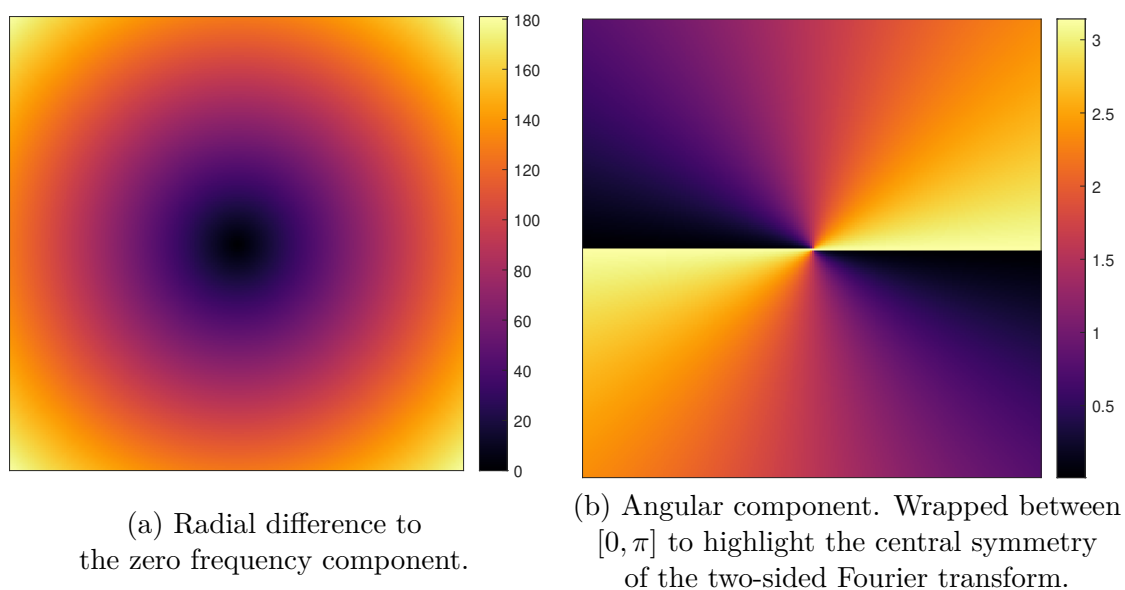


Figure 3-10: Examples of the visual representation of the polar coordinates of a power spectrum.

3-6. Radial binning with histogram equalization




Figure 3-11: Pipeline blocks relevant to radial histogram equalization and binning.

The reliability of statistical outlier labeling techniques, such as the ones in Sections 2-4-1-1 and 2-4-1-2 can be very sensitive to the sample size and consequently the sample’s ability to describe the underlying distribution adequately. The SOTA algorithms [70], [72], [74], described in Section 2-6, all use radial patterns in the Fourier representation as a foundation of their outlier labeling. These algorithms also have rather strong assumptions about the expected radial pattern, e.g. linearity on a log-log scale, or monotonously decreasing values as the radial distance from the zero frequency coefficient grows. Our approach attempts to use the radial information, without imposing strong assumptions on the radial pattern and without compromising the reliability of the statistical outlier labeling. We do this by assigning the Fourier coefficients to radial bins, consisting of a single circular bin in the center and further bins as concentric disks around it, this retains local radial information within the bins. However, in order to preserve the reliability of the statistical outlier labeling within these bins in a data-driven manner, we have to ensure that the sample size within the bins is controlled for images of arbitrary size. Histogram equalization provides a solution to this.

Histogram equalization is useful when most of the information within the image is contained in intensity values close to each other. Such images have poor contrast, and histogram equalization can help raise the contrast of those areas. This is done by reassigning frequently occurring intensity values to less frequently occurring values, such that the histogram of intensity values is as flat and uniform as possible across the range of possible intensities. This effect is shown on Figure 3-12. Figure 3-12a shows an image where all the intensity values are between $[0.45, 0.55]$, while the range of possible intensity values is $[0, 1]$. Figure 3-12b shows the histogram equalized intensities, spread out as uniformly as possible across the range of possible intensities. This raises contrast significantly. Mathematically this translates to the linearization of the cumulative distribution function, or in the discrete case, the empirical distribution function. In terms of radial binning, this increases the number of observations for bins with a low radial distance, and decreases the number of observations for bins with a high radial distance, resulting in a close to even distribution of observations across bins.

In the continuous case, in order to equalize the number of observations across radial bins, as a circle and disks, the goal is to designate concentric disks with increasing radii that cover an area equal to the initial circle. With discrete values, such as pixels of digital images and the Fourier transforms, a possible solution to this problem is radial binning using histogram equalization.

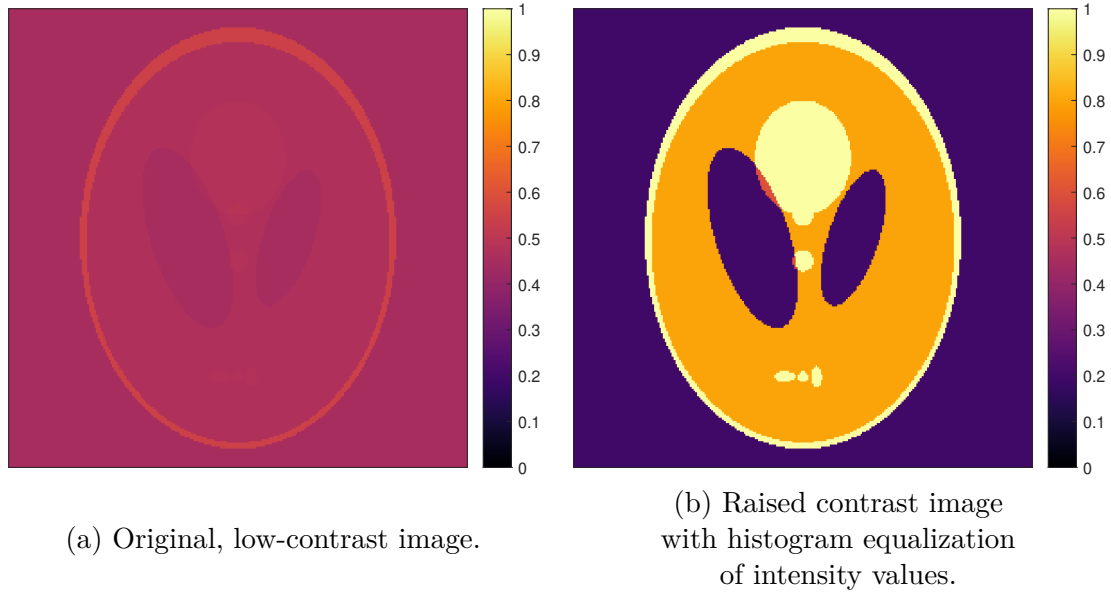



Figure 3-12: An example of increasing the contrast of an image through histogram equalization of intensity values.

Examples of equidistant and histogram equalized radial binning are shown in Figure 3-13. With equidistant binning, visible on Figure 3-13a, bins close to the center contain a low number of components, while bins far from the center contain an increasingly higher number of components. The histogram equalized binning strategy, visible on Figure 3-13b, contains an approximately equal number of coefficient values in each bin. This consistency in the number of components within each bin is beneficial to the reliability of the statistical approach for outlier detection, detailed in Section 3-9. Not having enough bins prohibits accurate approximation of the expected values at a certain radial distance that can make outlier detection unreliable. Having an unnecessarily large number of bins may result in inaccurate aggregates within bins, see in Section 3-7, where the bins do not generalize the underlying population, which in turn can give rise to an increased number of false positives in the outlier labeling process.



(a) Equidistant radial binning.

(b) Histogram equalized radial binning.

Figure 3-13: Examples of equidistant and histogram equalized radial binning from the central coefficient. With equidistant binning, bins close to the center contain a low number of coefficients, while bins farther from the center contain an increasingly higher number of coefficients. The histogram equalized binning strategy, contains an approximately equal number of coefficient values in each bin. Both figures show a zoomed-in, central section for visual clarity. The color coding indicates the index of the radial bin that the coefficient gets sorted into.

This is the first step where the pipelines for QID-RPCA and QID-MADN diverge. For QID-RPCA a masking has to be applied in order to execute the next step, described in Section 3-7. The radial bandwidth is limited for QID-RPCA because mapping polar values to a rectangular grid constrains the algorithm to the largest common radius for the 2-D Fourier transformed coefficients, which is directly proportional to the shorter spatial dimension. If this constraint is not enforced, the number of angular bins for a specific radial distance decreases, until there are only a handful of bins near the edge of the Fourier coefficient matrix. In turn, the decreasing number of angular bins make the outlier labeling strategy inaccurate and unreliable as it would be applied to a wildly different sample size, potentially resulting in false positive outliers. This phenomenon is not acceptable to our data-driven approach, thus only the Fourier coefficients within the largest common radius are used, while the coefficients outside this limit are excluded. This limit is equal to 128 components from the center, in the case of a 256 by 256 image. QID-MADN is not constrained by this common radius, as it does not need to transform to the rectangular representation of polar values. The effect of this is shown on Figure 3-14 for binning into 24 distinct bins, with and without the common radius masking.

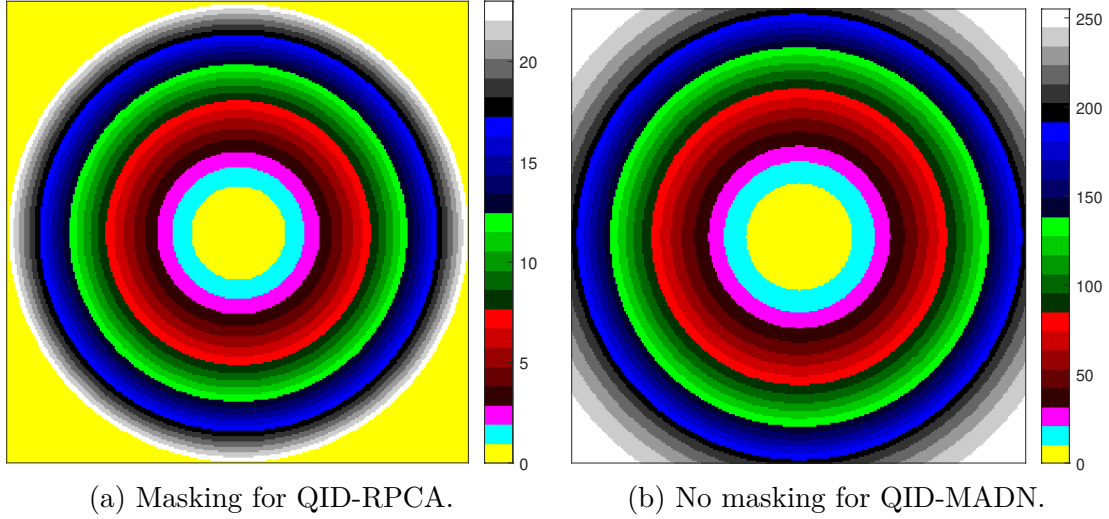


Figure 3-14: Examples of common-radius masking used for QID-RPCA and omitted for QID-MADN. Both histogram equalized binning strategies consist of 24 distinct bins.

The formulation of histogram equalization is described with Equations 3-7 through 3-13. N denotes the number of distinct radial bins used. The goal is to formulate a transformation function $T(r) = n$, where $r \in \mathbb{R}$ is the radial distance of the coefficient from the center and $n \in [1, N]$ is the index of the radial bin that the coefficient gets sorted into. The probability of a coefficient having a certain radial distance, r , from the center in the original setting is $p_1(r)$.

$$p_1(r) = \frac{\# \text{ of coefficients with radial distance } r}{\# \text{ of all coefficients}} \quad (3-7)$$

The probability of being a member of a certain radial bin, n , is $p_2(n)$.

$$p_2(n) = \frac{\# \text{ of coefficients in radial bin } n}{\# \text{ of all coefficients}} \quad (3-8)$$

Demonstrated for the continuous case. The relationship between $p_1(r)$ and $p_2(n)$ is as follows.

$$p_1(r) \frac{dr}{dn} = p_2(n) \quad (3-9)$$

Define the transformation function that maps a radial distance r to a radial bin n , as $n = T(r)$.

$$p_1(r) \frac{dr}{dT(r)} = p_2(T(r)) \quad (3-10)$$

The transformation function defined for histogram equalization is as follows.

$$T(r) = N \int_0^r p_1(q) dq \quad (3-11)$$

The derivative with respect to r .

$$\frac{dT(r)}{dr} = \frac{d(N \int_0^r p_1(q) dq)}{dr} = \quad (3-12a)$$

$$= N \frac{d(\int_0^r p_1(q) dq)}{dr} = \quad (3-12b)$$

$$= N p_1(r) \quad (3-12c)$$

Plugging this into Equation 3-9.

$$p_2(n) = p_1(r) \frac{dr}{dn} = \quad (3-13a)$$

$$= p_1(r) \frac{dr}{dT(r)} = \quad (3-13b)$$

$$= p_1(r) \frac{1}{N p_1(r)} = \quad (3-13c)$$

$$= \frac{1}{N} \quad (3-13d)$$

In the discrete case we can approximate the integral using a summation. The notation $\lceil \cdot \rceil$ corresponds to the ceiling function, which rounds values up to the closest integer.

$$T(r) = \lceil N \sum_{w=0}^r p_1(w) \rceil \quad (3-14)$$

The implementation calculates the rounded, expected number of coefficients in each bin based on the number of bins, N . Then sorts the array of coefficients according to radial distance and uses a pointer to perform the bin labeling on each of the sorted coefficients, until the required number of elements per bin has been reached. Due to the discrete nature of this task and because the number of bins is not required to be a divisor of the number of coefficients, the last bin can have slightly less or more coefficients than the rest. This error can be cancelled by constraining the number of bins to be a divisor of the number of coefficients. However, in practice, the error is marginal in most cases and does not affect the reliability of the outlier labeling. Thus, the constraint was deemed unnecessary for our implementations.

3-7. Mapping polar values to a rectangular grid




Figure 3-15: Pipeline blocks relevant to mapping polar values to a rectangular grid.

Note that this step is only done in the QID-RPCA variant. The next step for QID-MADN is found in Section 3-9. The QID-RPCA algorithm uses RPCA to separate the

Fourier coefficient matrix into two separate matrices, the low-rank matrix, which provides an expected value for the noise-free value and a sparse matrix that is expected to capture the components responsible for the quasiperiodic corruption, this is detailed in Section 3-8. Empirical experiments suggest, that using RPCA on the Fourier coefficients in the original coordinate system does not reliably separate the matrix into an accurate low-rank and sparse matrix. The goal of mapping polar values to a rectangular grid, is to highlight the radial pattern within the data and ensure that the matrix is separated into an accurate low-rank and sparse representation. Once the mapping to a polar coordinate system is complete and the radial values have been histogram equalized to the chosen number of radial bins, the matrix containing the Fourier coefficients is masked using each combination of radial and angular bins. For each set of values within that bin, the mean of the power is taken as the aggregate value of that bin in the rectangular grid. Figure 3-16 illustrates this combination of radial and angular bins. Looking at Figure 3-16c, the coefficients in each uniformly colored area are aggregated to one value in the rectangular grid, such that each radial disk transforms to a column in the rectangular grid and each angular slice transforms to a row in the rectangular grid. An example of this transformation of the power spectrum is visible on Figure 3-17, the rows correspond to a specific angular direction and the columns represent radial distances.

With a sufficient number of observations in every bin and a sufficient number of bins, it provides an approximate probability distribution of values for that radial distance. Using the mean value of the coefficients in a specific bin promotes the importance of high values in the rectangular representation, such that outliers can still be captured even if the majority of the coefficients in the original coordinate system are relatively small. This in turn makes it possible to label suspected outliers of that distribution. These outliers have a high probability that they are not part of the original distribution. Hence, they are unlikely to contain biological information and are likely caused by an external source of noise. Noise components are expected to affect only a handful of bins, meaning that most of the angular bins for a specific radius will have a low mean value, which is taken advantage of by RPCA, detailed in Section 3-8.

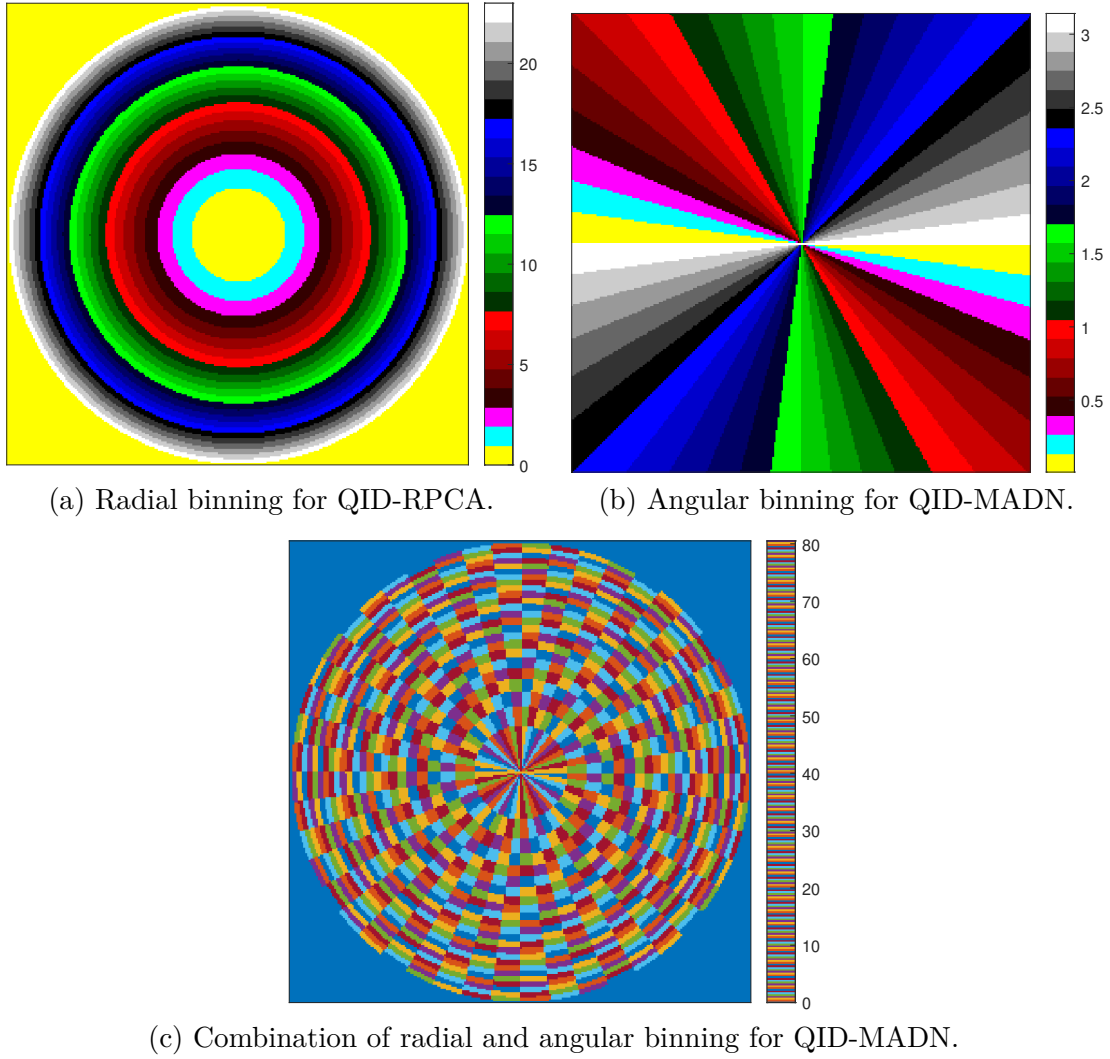


Figure 3-16: Example of combining radial and angular binning for the rectangular representation of polar values.

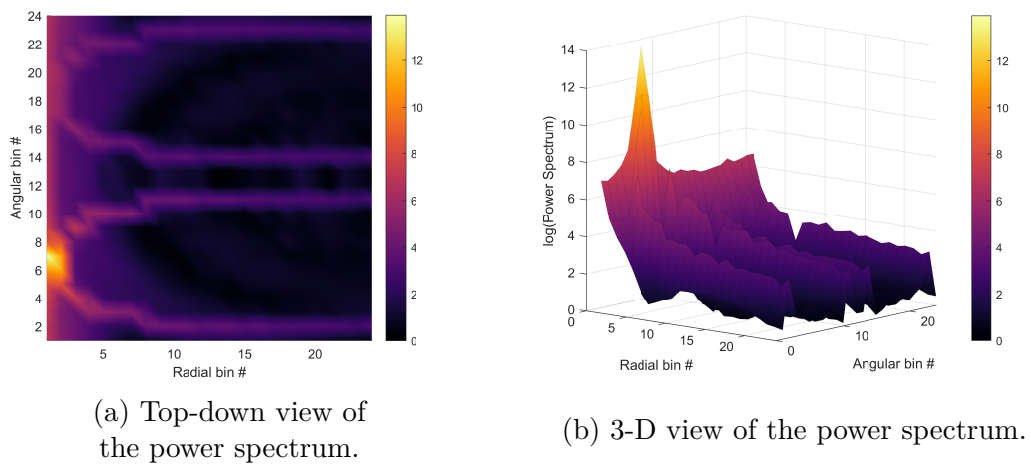


Figure 3-17: Example of the logarithm of the power spectrum, aggregated to the rectangular grid of polar values.

3-8. Robust Principal Component Analysis




Figure 3-18: Pipeline blocks relevant to robust principal component analysis.

Note that this step is only done in the QID-RPCA variant. The next step for QID-MADN is found in Section 3-9. RPCA, published by Candes et al. [58] in 2009, is a modification to Principal Component Analysis, which aims to separate the matrix into a low-rank component and a sparse component in the form of $M = L + S$, where $M \in \mathbb{R}^{n_1 \times n_2}$. This modification and its ability to separate out sparse variation into a separate matrix, S , from the low-rank structured matrix, L , makes it applicable to heavily corrupted data, which PCA would be sensitive to and would fit some of the first components to noise. The low-rank component, L , provides an expected value for that radial component, while the corrupt components are captured in the sparse matrix, S . The optimization problem is formulated according to Equation 3-15a.

$$\text{minimize} \quad \|L\|_* + \lambda \|S\|_1 \quad (3-15a)$$

$$\text{subject to} \quad L + S = M. \quad (3-15b)$$

$\|M\|_* := \sum_i \sigma_i(M)$ denotes the nuclear norm of M , which means the sum of the singular values of M , where the i -th singular value is $\sigma_i(M)$. $\|M\|_1 := \sum_{i,j} |M_{i,j}|$ denotes the l_1 vector norm, in other words, the sum of absolute values. According to the recommendation by Candes et al. the tuning parameter is set as described in Equation 3-16.

$$\lambda = \frac{1}{\sqrt{\max(n_1, n_2)}} \quad (3-16)$$

An example of this separation is seen on Figure 3-19.




Figure 3-19: Examples of the input and outputs of RPCA. The logarithm of the power spectrum in a rectangular, polar coordinate matrix M is separated into its low-rank component, L , and sparse component, S .

3-9. Outlier detection




Figure 3-20: Pipeline blocks relevant to outlier labeling.

Based on literature [42], one of the most commonly used approaches to detecting extreme observations, which separate well from the general patterns of the data, is through the use of dispersion and central tendency measures. The most commonly used dispersion measure is the standard deviation and the most commonly central tendency measure is the sample mean, often used for outlier detection and classification through the z -score. More details on the z -score are found in Section 2-4-1-1. Using this score, an observation is classified as an outlier, if its absolute difference to the sample mean is larger than n times the standard deviation of the sample, with $n = 3$ used by default. Even though this heuristic is simple, it is effective in many cases. Its robustness, however, is questionable, as the outliers themselves affect the mean and standard deviation of the sample. This

can result in a notably different detection threshold by only changing a few observations, while the distribution remains generally unaffected. To mitigate this effect, more robust measures of dispersion and central tendency can be used. A simple, yet robust extension to the three-sigma rule is through the use of normalized median absolute deviation (MADN) instead of the standard deviation and the median instead of the mean. More detail on the modified z -score is found in Section 2-4-1-2.

For QID-RPCA, outlier detection is done through calculating the z -score for every element of the sparse matrix S (notation taken from Equation 3-15a) and labeling elements that are over a z -score of three. Usage of the modified z -score is not possible in most cases as the sparse matrix has low cardinality. Low cardinality means that it contains a low number of unique values. In the current setting, this means that a few elements will take high values, while the majority of elements will be zero. The modified z -score has a breakdown point at the 50th percentile. In cases when at least 50% of the values are zeros within S , the median and MADN will be equal to zero. This is also called the exact fit property, which is not specific to the MADN, but happens to all robust estimators of scale [88]. When the exact fit property is satisfied, all elements of S that are non-zero are marked as outliers. This behavior is unacceptable in light of the denoising goal of minimal introduction of artificial patterns. Consequently, for QID-RPCA, the outlier detection is carried out through Z -scoring the elements of S , the sparse matrix and labeling elements over the z -score of three as outliers. Once outlier bins are identified, inversion of the rectangular, polar grid is carried out, creating a mask of outliers in the original, Cartesian coordinate system.

For QID-MADN, the rectangular binning step is omitted. This also prohibits the efficient use of RPCA as described before. However, robust outlier detection is still possible on the radially histogram equalized disks, by using the MADN outlier labeling approach on each component individually. The median of a disk is used as the local, expected value for the disk, while classifying individual coefficients as outliers if their distance to this median is larger than three times the normalized median of absolute differences of that specific disk. This is described in Equation 3-17, where $PS_{k,l}$ is the power coefficient at indices k and l in the Cartesian coordinate system and d is the index of the disk, which $PS_{k,l}$ is contained in. The function $MADN()$ is described in Equation 2-15. Apart from computational simplicity, this approach is advantageous compared to the rectangular binning approach as components are evaluated individually. This way, no change occurs in a large number of components in the region of the outliers, while with the rectangular binning approach of QID-RPCA, every component in a specific bin is altered. Leakage of these discrete coefficients is still possible, since the resolution of the Fourier representation is limited. Therefore each component that is identified as an outlier in QID-MADN is dilated to adjacent components in the outlier map, with a default radius of three coefficients. This counteracts most of the residuals caused by leakage, while also being less invasive compared to the QID-RPCA binning approach.

Examples of both filtering approaches are illustrated in Figure 3-21. The outlier map of QID-RPCA is shown both in the rectangular, polar coordinate-system, as well as the original Cartesian coordinate-system.

$$\log(PS_{k,l}) - \text{median}(\log(PS_d)) > 3 * MADN(\log(PD_d)) \quad (3-17)$$

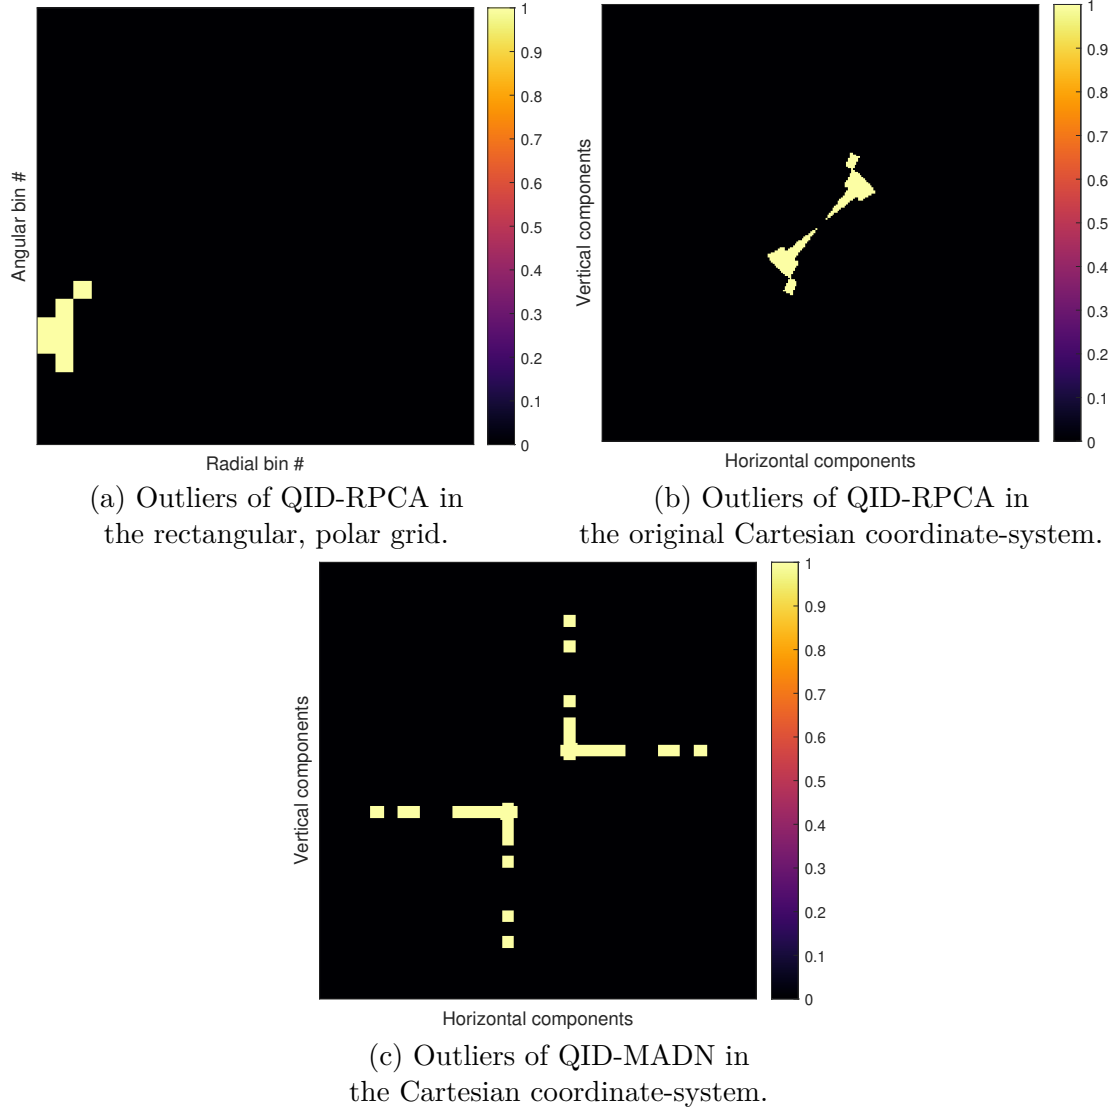


Figure 3-21: Examples of the outlier maps of QID-RPCA and QID-MADN. The outlier map of QID-RPCA is shown both in the rectangular, polar coordinate-system, as well as the original Cartesian coordinate-system.

3-10. Optimization of the binning strategy

As mentioned earlier in Sections 3-6 and 3-7, too few, as well as too many bins may affect the performance of QID-RPCA significantly. Determining the number of radial and angular bins that ensures the accurate labeling of outliers in a data-driven manner before evaluating the denoised image, however, is not trivial. Until this point, no explicit criteria have been established regarding the optimality of the binning strategy. The goal is to find a binning strategy, which provides maximal separation between the low-rank and sparse component in RPCA, without the loss of generality and with the minimal computational costs possible. The sparse component, S , serves as the basis for outlier labeling, while the low-rank component, L , serves as the expected, noise-free value of a specific component. Hence, optimality of the binning strategy should be determined based on the sparsity of S , as it directly impacts the detection of outliers. The automated binning strategy for QID-MADN is found at the end of this section.

Minimal sparsity in any sample, such as the values of the sparse matrix, S , occurs when

the values are uniformly distributed within the sample, while it is maximally sparse when all elements are zero valued, except for one. This definition is not sufficient, it does not control for a large number of variables, e.g. sample size. As a result, some measures of sparsity fail to provide a precise and straightforward comparison in some instances. In the financial field, measuring the distribution of wealth and its inequality on different scales, across different neighborhoods, countries, and continents, is often done through sparsity. Much of the analysis at hand stems from learnings in the financial sector, such as Dalton's laws [4]. Terminology such as 'rich' and 'poor', 'wealth' and 'energy' is also borrowed for a more intuitive description of the phenomena, even though it has no direct relevance to medical imaging. Hurley and Rickard [52] provide a detailed comparison of sparsity measures, highlighting their strengths and weaknesses through a set of criteria. The criteria used by Hurley et al. to determine the quality of a sparsity measure include the following desirable characteristics. Choosing a sparsity metric that satisfies as many of these characteristics as possible is also desirable in our case, as it may predict the outlier detection performance of QID-RPCA.

- Robin Hood: taking from the 'rich', coefficients with a high value, and giving to the 'poor', coefficients with a low value, while introducing no wealth or energy to the population, decreases sparsity.
- Scale invariance: the sparsity measure should not be sensitive to the energy contained in the system. Absolute values of the population do not matter, only the relative differences between the minimally and maximally sparse case.
- Rising tide: adding a constant to the energy of every member of the population decreases sparsity.
- Cloning: the sparsity measure is invariant to population size. Two populations with identical sparsity, when combined, result in the same sparsity.
- Bill Gates: as one member of the population increases in energy infinitely, while the others stay constant, the sparsity measure asymptotically grows towards its maximum value.
- Babies: adding members with zero wealth increases sparsity.

Their analysis covers a large number of vector norms, such as the absolute-value norm (L_1) or the maximum norm (L_∞), as well as metrics such as the Hoyer [37] and the Gini index [6]. Many of the measures analysed by Hurley et al. are inherently dependent on and scaled by sample size. The number of radial and angular bins, determines the number of elements in the sparse matrix, S , hence the sample size. This is especially critical for us, as measuring the sparsity of different binning strategies has to account for the fact that these strategies compare different sample sizes. Providing simple, but intuitive descriptions of the norms in the context of the application may aid finding a reasonable measure. The absolute-value norm, also known as the L_1 norm, when applied to the power spectrum of an image, is proportional to the total power captured by the signal. The higher the total power, the higher the norm. The Euclidean-norm, or L_2 norm, on the other hand, apart from the total power is also sensitive to the concentration of power in a low number of observations. Two vectors with the same total power will have a higher L_2 norm if the power is concentrated in a single observation or coefficient, compared to if it is distributed uniformly. The ratio between these norms, $\frac{L_2}{L_1}$, indicates how concentrated a single unit of power is. When all the power is concentrated in a single observation, this measure is equal to 1, independent of the population size. While it is monotonically decreasing until the power is uniformly distributed among all the observations, in which case the value is $\frac{1}{\sqrt{N}}$, where N is the population size. This measure satisfies the Robin Hood, scale invariance,

and Bill Gates criteria, while violates the rest [52].

Hoyer utilizes min-max scaling on $\frac{L_2}{L_1}$ to map it between $[0,1]$ independent of population size, shown in Equation 3-18, where N is the number of elements in the matrix, L_1 and L_2 are the norms of the same matrix. With this modification, the sparsity measure satisfies all criteria except for Cloning [52], while still indicating how concentrated a single unit of power is. Because of its favorable properties and intuitive interpretation, the Hoyer metric is used by us to choose the optimal binning strategy, based on the sparsity of S in QID-RPCA.

$$Hoyer = \frac{\sqrt{N} - \frac{L_1}{L_2}}{\sqrt{N} - 1} \quad (3-18)$$

A range of angular and radial bin combinations are used for binning to the aggregate values of a polar bin to the rectangular grid, next RPCA is used, then the Hoyer sparsity metric is calculated for each sparse matrix, S . The binning strategy resulting in the highest Hoyer metric is chosen. Since the sparse matrix of RPCA is used for determining optimality, calculating it is necessary for every binning strategy. This can be computationally intensive. In most practical cases, the optimal strategy comes out to be approximately $\sqrt{\min(width, height)}$ for both the number of radial and the number of angular bins. For applications such as multispectral immunofluorescent microscopy, where the same biological sample is imaged multiple times, this optimality calculation only has to be done once and the results can be used for each separate band of the sample.

For QID-MADN, denoising performance is less sensitive to the composition of the bins. As opposed to QID-RPCA, no aggregation and no transformation to a rectangular polar grid is done. Furthermore, MADN filtering is applicable to relatively low sample sizes and has a breakdown point of approximately 50%, which means that it is not sensitive to outliers [14]. Through empirical testing, the number of bins can robustly be set to \sqrt{n} , where n is the number of pixels. While a minimum sample size of 30 measurements is commonly recommended in order for a sample to approximate a distribution, determining the minimum sample size depends on a number of statistical parameters and assumptions [3]. Thus, the minimum sample size should be determined through sample-specific, statistical power analysis. The binning strategy for QID-MADN is summarized in Equation 3-19, where N is the number of bins and n is the number of pixels in the original image. A lower bound of 64 coefficients per disk is set, in order to steer clear of relatively low sample sizes and due to the unlikely nature of biomedical images under the resolution 64×64 . An upper bound of 1024 disks is set, as for large images the computational cost of using more bins out-weighs the potential increase in resolution and precision in attenuation.

$$N = \begin{cases} 64, & n < 4096 \\ \sqrt{n}, & 4096 \leq n < 262144 \\ 512, & 262144 \leq n \end{cases} \quad (3-19)$$

3-11. Noise component attenuation






Figure 3-22: Pipeline blocks relevant to noise component attenuation.


Bins and elements which are classified as outliers by QID-RPCA and QID-MADN, are attenuated to their best estimate of their noise-free value. This best estimate of what value the labeled coefficients should take if no noise were present for QID-RPCA is captured by the low-rank component of RPCA, L . The best estimate for QID-MADN is the median of the respective radial disk. The advantages of noise component attenuation compared to noise filtering is explained in more detail in Section 2-4-2-3. Examples of attenuated power spectra for each algorithm are shown on Figure 3-23. Figure 3-21b and 3-21c show the labeled outliers attenuated in Figure 3-23.



(a) Power spectrum of the original, corrupt image.



(b) Attenuation by QID-RPCA.



(c) Attenuation by QID-MADN.

Figure 3-23: Comparison of the original and attenuated power spectrums.

3-12. Image reconstruction by inverse Fourier transform





Figure 3-24: Pipeline blocks relevant to transform inversion.

The final step of the pipeline is image reconstruction through the inverse Fourier transform, resulting in a denoised image in the original domain. The complex valued Fourier representation of the filtered image is retrieved through Equations 3-20 and 3-21, where $P_{p,r}$ is the power of a specific coefficient, $\Phi_{p,r}$ is the phase of the same coefficient, and the phase is retained from Section 3-3. For an $n \times m$ matrix in the spectral domain, Y , and its inverse Fourier transform X in the spatial domain. $\omega_m = e^{-\frac{2\pi i}{m}}$ and $\omega_n = e^{-\frac{2\pi i}{n}}$ are complex roots of unity, where i is the imaginary unit, p and q are indices in the Fourier domain, j and k are indices in the spatial domain, p and j are indices that run from 0 to $m-1$, and q and k are indices that run from 0 to $n-1$. An example of the results of this inverse transform is seen on Figure 3-25.

This step concludes the methods used in QID. For a comprehensive overview, please refer back to Section 3-1 and Figure 3-2.

$$Y_{p,r} = \sqrt{P_{p,r}} e^{i\Phi_{p,r}} \quad (3-20)$$


$$X_{j,k} = \frac{1}{m} \sum_{p=0}^{m-1} \frac{1}{n} \sum_{q=0}^{n-1} \omega_m^{pj} \omega_n^{qk} Y_{p,q} \quad (3-21)$$



(a) Original, corrupt image.



(b) Denoising by QID-RPCA.



(c) Denoising by QID-MADN.

Figure 3-25: Comparison of the original and denoised images.

Chapter 4

Case studies

In order to demonstrate the applicability and usefulness of QID and compare its performance to state-of-the-art algorithms, a series of case studies are carried out. This section is organised around three sets of experiments on three separate datasets. Namely, a synthetic dataset is used to paint a broad picture of the properties and specific behaviors exhibited by the algorithms and two, real-world datasets are used to compare the approaches in real-world scenarios. First, in each case study, the datasets are introduced, background information on their origin is given, and motivation behind their use is described. Next, the experiments are outlined, but also what the expected results would be and possible shortcomings are.

The datasets used are as follows:

- a synthetic dataset based on the Shepp-Logan phantom [10],
- a set of processed images based on a brightfield microscopy image of a transversal section of a mouse brain and a coronal section of a rat kidney,
- a set of raw, fluorescence microscopy images of a resected human kidney.

4-1. Shepp-Logan synthetic dataset

The first set of experiments use a series of artificial noise patterns, with varying parameters. As the base image the Shepp-Logan phantom [10] is used, seen on Figure 4-1. It is an artificial image that models a human head section by x or gamma radiation imaging [11]. It is created in order to help the development of Fourier reconstruction algorithms medical images. It simulates different biological structures, e.g. skull, gray matter, tumors, and hematoma.




Figure 4-1: Modified Shepp-Logan phantom.

Choosing the established Shepp-Logan phantom has advantages. In order to properly approximate real performance in the medical imaging domain, the image used cannot be chosen in an arbitrary fashion, it has to possess features that are similar to those in the medical imaging domain. Images that are too simplistic, such as a classic checkerboard pattern include long, straight lines of sharp transitions. While this pattern would e.g. amplify the severity of Gibbs-phenomena, such a pattern would occur infrequently in medical images. Specific artificial patterns that are not found in real medical images may hide important algorithmic performance issues that only become visible in more complex cases. On the other end of the spectrum, complex natural images, commonly used for testing of signal and image processing algorithms such as "The mandrill", "peppers", "Lena", or "The cameraman" miss some aspects and features specific to the medical imaging domain. Specific biological structures, such as bones, cell membranes, and tissue textures follow patterns that are not appropriately modelled by such images. Furthermore, these natural images do not contain buffer areas with no useful information, such as masks around viewports or empty viewports outside of the biological sample. Using natural images as a basis of our synthetic experiments would risk focusing on issues not found in the medical domain.

4-1-1. Experimental setup

Determining the denoising performance for quasiperiodic corruption accurately and reliably is a complicated task by itself. In the ideal case, one knows the exact noise models possible, with a full understanding of their parameters, such as amplitude, direction, and distribution. Furthermore, one would have a dataset of real images containing all of these corruptions, labeled, as well as the pure, noiseless counterparts for each. For medical im-

ages, this is rarely, if ever, available. In most cases, obtaining an objectively pure, noiseless sample is simply not possible.

With the experiments on this synthetic dataset, getting close to this ideal case is the goal. Creating labeled examples of corruptions, across a wide range of parameters used for structured comparison is essential. Furthermore, evaluation of specific phenomena, such as information loss through saturation and added noise over the Nyquist frequency is examined using this dataset for which the ground truth is available.

First, limits and behaviors of the algorithms are compared using this dataset employing a simple, deterministic corruption scheme. This way, it is easier to establish causal relationships between the change in a specific variable and its effects on denoising performance. For example, we use pure sinusoids, which are periodic functions, to determine performance on quasiperiodic corruptions, albeit a significant simplification of quasiperiodic corruption in the wild. These periodic patterns, are however sufficient to determine at least the upper and lower detection limits to corruptions in terms of frequency, as well as the amplitude. Even though the complexity is significantly reduced with this synthetic approach and certain variables present in the real-world examples are disregarded here, it does allow for analysis of specific phenomena in a rather structured manner, which is rarely a possibility with real datasets.

The experimental setup is as follows. Starting from the baseline image shown in Figure 4-1, a single sinusoidal pattern is added, with a specific frequency and amplitude combination, while phase is kept constant throughout. Frequencies are varied between 0.02 cycles/pixel and 0.5 cycles/pixel, in increments of 0.02 cycles/pixel. A corruption with a frequency of 0.02 cycles/pixel is low enough that it's unlikely to occur in biomedical images, and a frequency of 0.5 cycles/pixel is high enough to simulate the limit imposed by the Nyquist frequency. Amplitudes are varied between 5% and 25% of the available pixel intensity scale, with changes in increments of 5%. This results in a set of 125 distinct corrupted images for which the ground truth is known, since we introduced the noise. In order to counteract information loss through intensity saturation, the intensity range of the baseline image is shrunk to the central 50 percentile. This means that with an intensity range of $[0,1]$, the clean, ground truth image will contain values in the range $[0.25,0.75]$, while the corrupted version with a sinusoid with amplitude of 25% if the intensity and an arbitrary frequency will occupy the full intensity range at $[0,1]$.

Once the dataset is generated, a high-level comparison and performance analysis of denoising using state-of-the-art (SOTA) algorithms, namely ARPENOS, ACARPENOS, and IONITA, described in Section 2-6 and our QID methods are carried out. The evaluation of denoising performance makes use of a set of metrics, namely RMSE, SSIM, SNR, PSNR, and BRISQUE. Then, we show concrete examples, in the image domain to highlight specific artifacts and observations. The metrics listed above are sensitive to some artifacts, while oblivious to others, as described in Section 2-5-2. Pižurica et al. [45] further argues that for applications where the final result is to be used by humans, final evaluation by humans is often the best and only reliable way of evaluation.

Finally, this case study also includes special cases using modified datasets, with the aim to give an indication of how these algorithms behave in a wider range of real-world scenarios.

The list of special cases are as follows:

- how well they handle information loss through intensity saturation;
- how well they handle corruption patterns with frequencies over the Nyquist-frequency;

- and how well they scale in terms of computation time.

The algorithms compared include ARPENOS and ACARPENOS by Sur et al., the algorithm by Ionita et al., which is referred to as *IONITA* from now on, and out two approaches described in Chapter 3, QID-RPCA and QID-MADN.

4-2. Processed mouse brain and rat kidney dataset

This dataset contains a labeled set of processed images, that originate from single microscopy images of a transversal section of a mouse brain and a coronal section of a rat kidney. The processing includes different color spaces such as RGB, HVS, and YCbCr, specific factors of decomposition and factorization methods such as PCA, NNMF, and DWT, range- and median filtering, entropy- and median filtering. A set of examples can be seen in Figure 4-2 with a custom color scale. The dataset also contains combinations of these resulting in a vector of 905 distinct images with a varying level of noise and signal in each. All of them originate from a single microscopy image of a transversal section of a mouse brain and a coronal section of a rat kidney, respectively.




Figure 4-2: Processed microscopy images, on a false color scale. Processing includes different color spaces such as RGB, HVS, and YCbCr, specific factors of techniques such as PCA, NNMF, and DWT, range- and median filtered, entropy- and median filtered versions.

4-2-1. Experimental setup

In this case, there is no objective, noiseless image available. As a result, it is more difficult to give an unbiased comparison between the performance of the algorithms. However,

a previously introduced reference-free quality metric, BRISQUE, can still be used. Even though it will be biased towards a specific set of features, it provides a basis for comparison. Furthermore, image domain comparisons are shown to enable human interpretation of the denoising results, highlighting specific behaviors and artifacts.

4-3. Raw human kidney dataset


An instance of structured noise is a phenomenon called viewport-stitching. High-resolution microscopy images, such as this dataset for example, are not often acquired by a single high-resolution imaging sensor, at a single instant in time. They are often acquired by a lower-resolution sensor that takes an image of a small part of the sample, this is what we refer to as a viewport. This sensor is then then moved along a specific structured path, such that a lower-resolution image is taken of each and every part of the sample. When all the areas of the sample have been captured, the lower-resolution images or viewports are connected or stitched together, this results in a high-resolution image. The amount of light that falls on the sample within and across viewports is often not uniform, especially without regular calibration and maintenance. When stitching these non-uniformly lit viewports together, a structured, repetitive pattern emerges, which is referred to as viewport-stitching artifact or viewport corruption [27].

The immunofluorescence image data that is contained in this dataset was collected from a 10 μm thick cryo-section of a resected human kidney. A Zeiss Axio.Scan.Z1 slide scanner was used to capture the image using a 10x Plan-Aprochromat/0.45 NA M27 Zeiss objective and a Hamamatsu ORCA Flash 2.0 camera with 16-bit precision. A metal-halide Zeiss HXP-120 fluorescence lamp was used as light source with 90% intensity used for all channels. The image contains 4 immunolabels captured through 4 fluorescent filter channels described in Table 4-1.


Filter name	Filter set (Zeiss)	Item no. (Zeiss)	Filter excitation wavelengths (nm)	Filter emission wavelengths (nm)	Marker	Structure marked	Exposure time (ms)
DAPI	Filter Set 49	488049-9901-000	335-383	420-470	Hoescht	cell nuclei	20
EGFP	Filter Set 38 HE	489038-9901-000	450-490	500-550	Laminin	basement membrane	320
DsRed	Filter Set 43 HE	489043-9901-000	538-562	570-640	Synaptopodin	glomerular	89
Cy 5	Filter Set 50	488050-9901-000	625-655	665-715	THP	thick limb	250

Table 4-1: Table of filter and immunomarker information for the human kidney immunofluorescence image. Each row describes one of the four band examined in this experiment.

As an example, the raw EGFP band is shown in Figure 4-3. Figure 4-3b shows the contrast stretched image for increased visibility of the low-intensity patterns captured by this band. Note that over 50% of the pixels on this image do not contain biological information and are "off tissue". The majority of these pixels are not physically measured, located around the viewport blocks which are physically measured, to complete a rectangular image. A smaller number of pixels are physically measured without capturing biological signal, in the viewports at the edges of the kidney tissue and the biological information. These empty and near-empty viewports are necessary in most cases to guarantee that every area with biological tissue information is adequately captured. The mask of measured viewports is clearly visible in the uniform black regions of Figure 4-3, while the viewports without biological signal are easy to identify by their close to uniform dark purple color.



(a) Raw microscopy image.



(b) Saturated microscopy image for increased visibility.

Figure 4-3: EGFP band of the human kidney microscopy dataset. Note the grid-like repetitive pattern resulting from viewport being stitched together, superimposed on top of the biological pattern of the human kidney.

4-3-1. Experimental setup

Until now, all experiments have been conducted on either synthetic images with controlled corruptions or on real-world microscopy images processed in such a way that certain pat-

terns get amplified or attenuated. Even though, this provides valuable information on the strengths and weaknesses of the algorithms, it does not sufficiently showcase performance on the intended use-case, namely raw biomedical images. This third dataset, however, contains unaltered real-world microscopy images, exactly as they are acquired by the instrument. The algorithms are used without any preprocessing, knowledge of noise characteristics or parameter tuning, making application of the QID algorithms on this dataset a true test as envisioned at the outset of this thesis.

In this third case, there is no objective, noiseless image available. As a result, it is also more difficult to give an unbiased comparison between the performance of the algorithms. The previously used no-reference metric for natural images, BRISQUE, is expected to be ineffective for quality assessment for these images because the mask and empty viewports take up over half of the pixels in total, which render the test images far from the natural images that BRISQUE assumes. Since no information is given to the algorithms to ignore or handle these areas differently, modifications may be made to them, which in turn may influence the 'naturalness' of the images even further, even when the biological signal is denoised. As a result, no reliable BRISQUE score can be shown to evaluate these images. With no reference image ruling out RMSE etc. and no suitable no-reference metric available in BRISQUE, the focus in the algorithm evaluation on this third case study lies primarily on comparisons in the image domain with human inspection at the forefront.

Chapter 5

Results and Discussion

5-1. Synthetic dataset results and performance comparison

As explained in Section 4-1-1, starting from Figure 4-1, a single sinusoidal pattern is added to mimic a quasiperiodic corruption, each with a specific frequency and amplitude combination, while the phase is kept constant throughout. Frequencies are varied between 0.02 cycles/pixel and 0.5 cycles/pixel, in increments of 0.02 cycles/pixel. Amplitudes are varied between 5% and 25% of the available pixel intensity scale, with change occurring in increments of 5%. This results in a set of 125 corrupted images for which the ground truth is known.

Once the dataset is generated, a high-level comparison and performance analysis of denoising are carried out. The comparison is done by using SOTA algorithms, specifically ARPENOS, ACARPENOS, and IONITA and out newly developed approaches, QID-MADN and QID-RPCA. Parameters for ARPENOS, ACARPENOS, and IONITA are chosen according to the recommended values found in their respective papers [70], [72], [74]. In certain cases, some parameters, e.g. ones that scale with image size are adjusted to the image. Most parameters are tuned to reflect the best denoising performance possible. However, in some instances, a more fine-tuned set of parameters may exist, despite best efforts. Evaluation of performance is based on a set of metrics, namely RMSE, SSIM, SNR, PSNR, and BRISQUE. RMSE, SNR, and PSNR are metrics that are widely used in science and engineering, which quantify the signal strength and the presence of noise. SSIM on the other hand incorporates aspects of the human visual system into its scoring, and is more in-line with the human perception of quality. BRISQUE is used to test the usefulness and applicability of a no-reference metric, which may be the best tool for numerical quality assessment in the real-world datasets where no ground truth is available. After the numerical assessment, concrete examples are shown, in the native image domain in order to compare specific artifacts across the algorithms. Pižurica et al. [45] further argues that for applications where the final result is to be viewed by people, final evaluation by people is often the best and only reliable way of evaluation. A perceptually uniform color scale, *inferno*, has been used in order to promote visibility and uniformity across the different images.

5-1-1. Frequency domain performance comparison of QID-RPCA and QID-MADN versus prior algorithms from the literature

In this section, a high-level comparison and performance analysis of quasiperiodic denoising using the SOTA algorithms, specifically ARPENOS, ACARPENOS, and IONITA and the previously described novel approaches, QID-MADN and QID-RPCA are carried out. Evaluation of performance is based on a set of metrics, namely RMSE, SSIM, SNR, PSNR, and BRISQUE. This means that the full table of scores consists of 125×5 algorithms \times 5 metrics = 3125 values. Comparing these scores and their patterns comprehensively in a tabular format is impractical. Instead, absolute scores of QID-RPCA and QID-MADN are shown in a graphical visualization, using a grid of frequencies and amplitudes of the added corruption and evaluation scores indicated as color. Furthermore, comparison of the current SOTA to both QID algorithms is executed through a similar visual format, only the relative scores to QID-RPCA and QID-MADN are shown to clearly communicate where the differences lie for the different algorithms. For the absolute scores depicted in the QID visualizations, a lighter color means better denoising performance while a dark one means worse, irrespective of the metric. For relative scores between QID and one of the SOTA algorithms, the SOTA score is subtracted from QID. A divergent color scheme is used to highlight differences. Blue means that QID performs better, red means that the SOTA algorithm performs better, while white means, that both algorithms perform similarly. Note that a table of discrete quantitative metric scores could give a misleading impression of a precise arithmetic assessment, while, colors provide a wider spectrum that allows for a more intuitive and visual interpretation of the results exhibited by each algorithm.

In terms of mean squared error (MSE), all the algorithms perform well on high-frequency noise components. As seen on Figures 5-1 and 5-2, most algorithms have a definite and quite stark low-frequency cutoff point under which denoising performance declines significantly. For both QID algorithms, residuals and errors in denoising start to grow considerably under 0.1 cycles/pixel. However, relative to the others, the highest error is still approximately one order of magnitude lower, peaking at around 4.5×10^{-3} . For the SOTA algorithms this is 3.3×10^{-2} . It is important to note that for QID this limit is affected by the zero-frequency protection radius, which is a circle with a radius of 5 pixels by default, corresponding to approximately 4% of the full frequency range in this case. For IONITA, this low-frequency limit is around 0.12 cycles/pixel, for ACARPENOS it is 0.1 cycles/pixel, for ARPENOS the limit is 0.05 cycles/pixel.

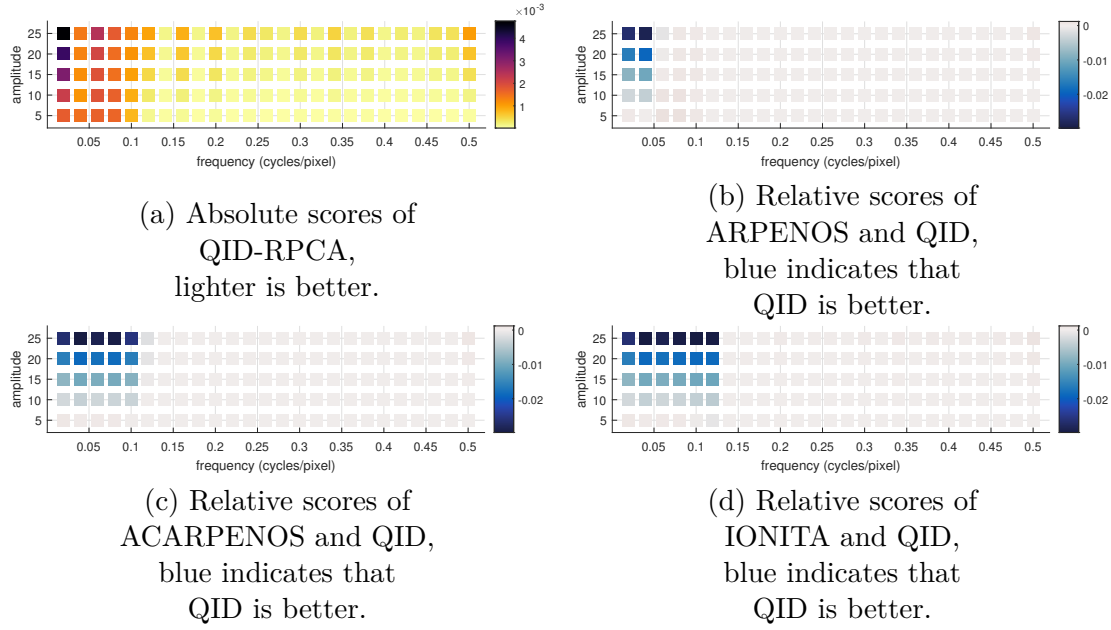


Figure 5-1: Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of MSE score.

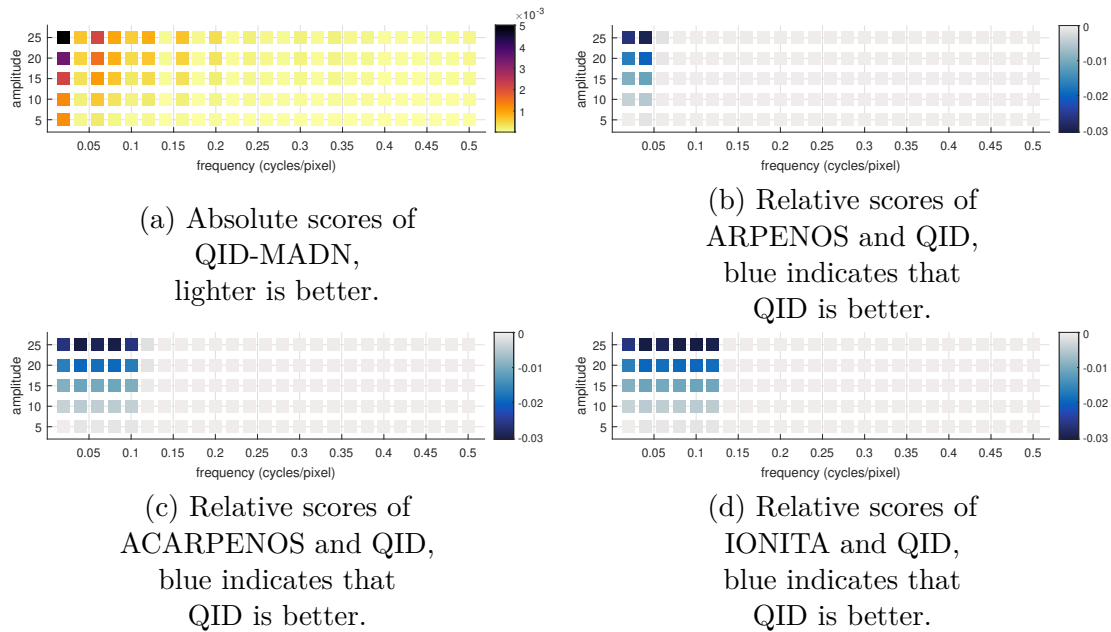


Figure 5-2: Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of MSE score.

Structural similarity (SSIM) focuses more on the perceptual quality of the denoised image, since it is more in line with the human visual system (HVS) compared to widely used error metrics such SNR and MSE. This is why SSIM is commonly used for tasks where aiding human interpretation is the goal, such as in the case in the current medical image denoising problem. This metric is used as the primary metric here because perceptual quality of the image is of the highest importance.

As seen on Figures 5-3 and 5-4, the aforementioned low-frequency limit is still visible. However, its boundaries shift for this metric. In the case of ACARPENOS and IONITA, the

lower bound for successful denoising seems to be between 0.12 and 0.14, while ARPENOS stays between 0.04 and 0.06 cycles/pixel, . Under this bound, both QID algorithms perform significantly better than the algorithms in literature. Over this bound, performance of the QID and SOTA algorithms is similar, with minor differences.

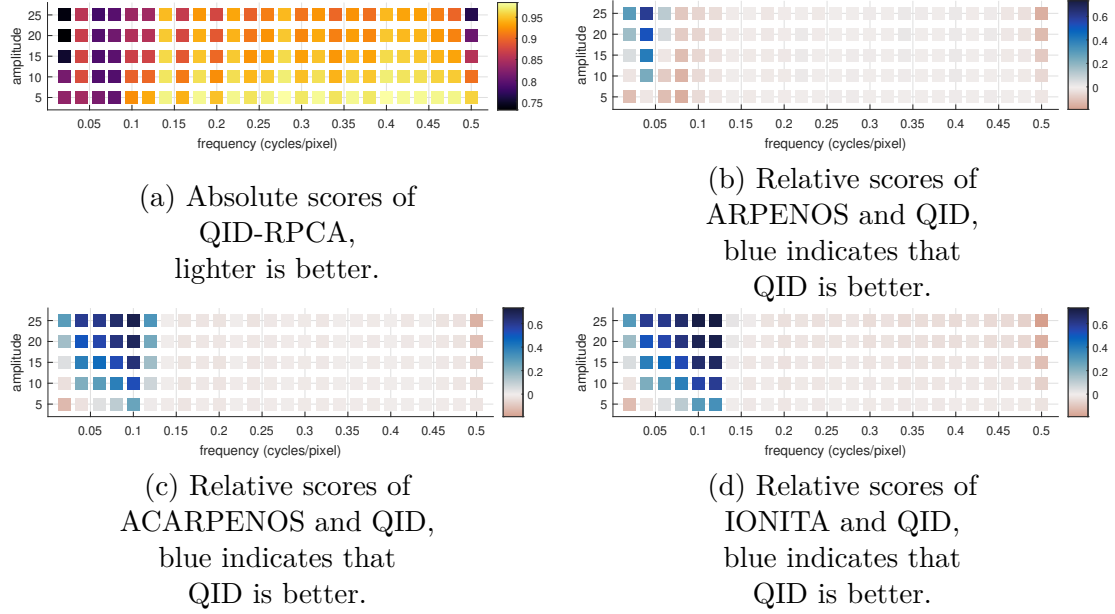


Figure 5-3: Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of SSIM score.

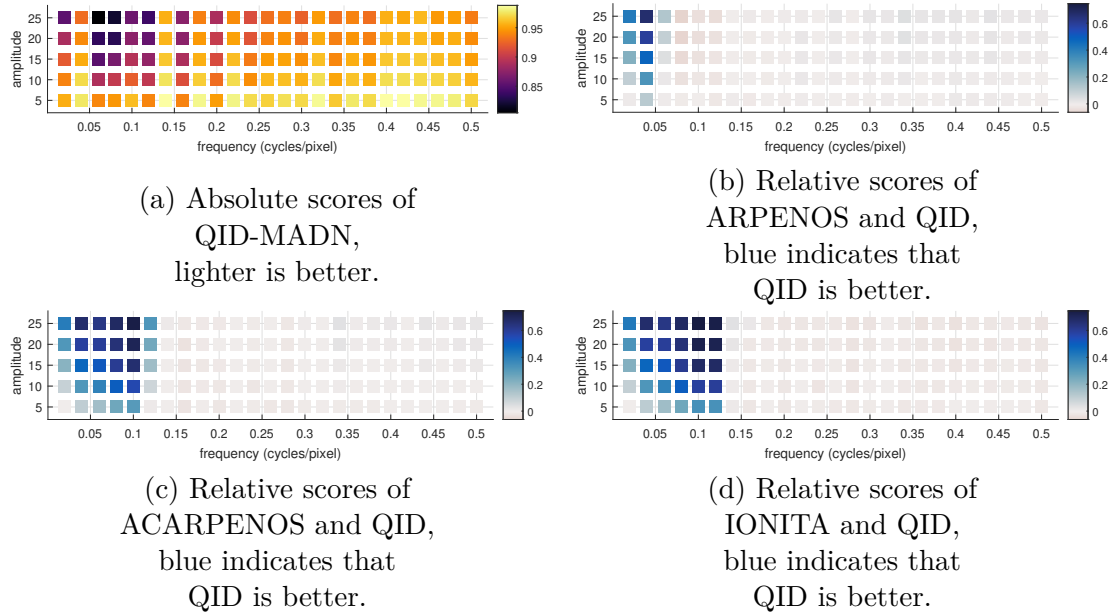


Figure 5-4: Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of SSIM score.

Signal-to-noise ratio (SNR) is a widely used metric within the signal processing community. SNR tends to highlight more minute differences compared to MSE and SSIM. As seen on Figures 5-5 and 5-6, a low-frequency boundary is still visible between 0.12 and 0.14 for ACARPENOS and IONITA, while the stable 0.04 to 0.06 boundary is observable for ARPENOS. QID-RPCA maintains an approximately 3 dB higher score for

high-frequency corruptions, except around 0.5 cycles/pixel, near the Nyquist frequency, where performance declines. For high-frequency noise components ARPENOS is consistently worse than QID-MADN by approximately 5-6 dB, the difference grows to 15 dB in the low-amplitude range. In terms of SNR, IONITA outperforms QID-RPCA in the high frequency range. The difference continuously grows for higher frequencies, and near the Nyquist frequency it reaches a difference of approximately 6 dB. QID-MADN on the other hand only performs marginally worse than IONITA, the highest difference being approximately 3 dB, showing the first real instance of QID-MADN outperforming QID-RPCA. ACARPENOS performs worse than IONITA and better than ARPENOS, but shows no particularly interesting behavior compared to its alternatives.

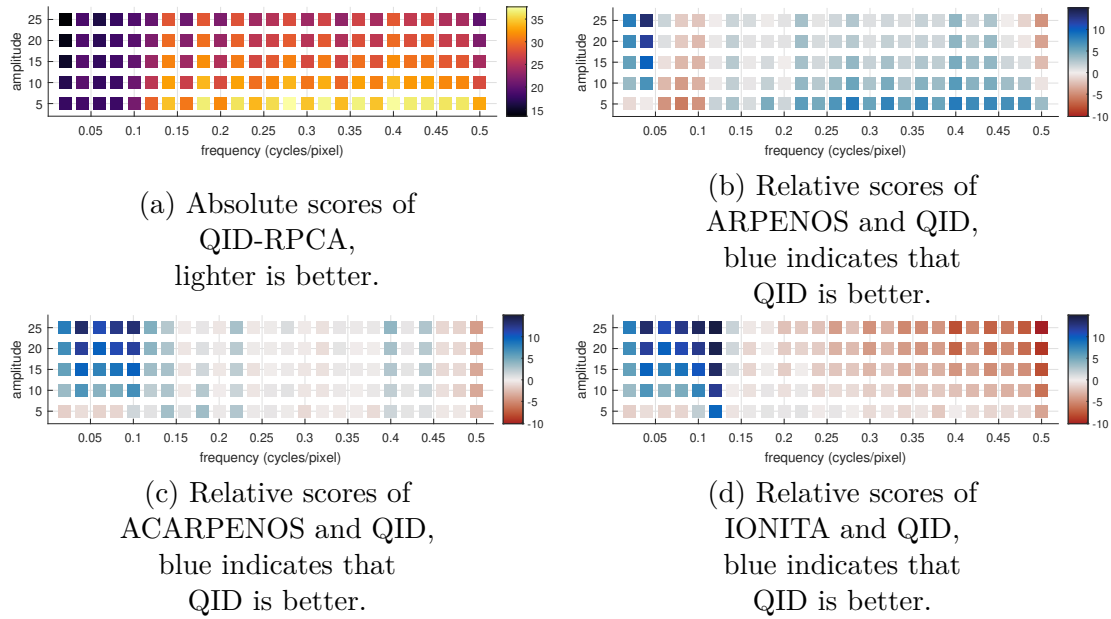


Figure 5-5: Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of SNR score.

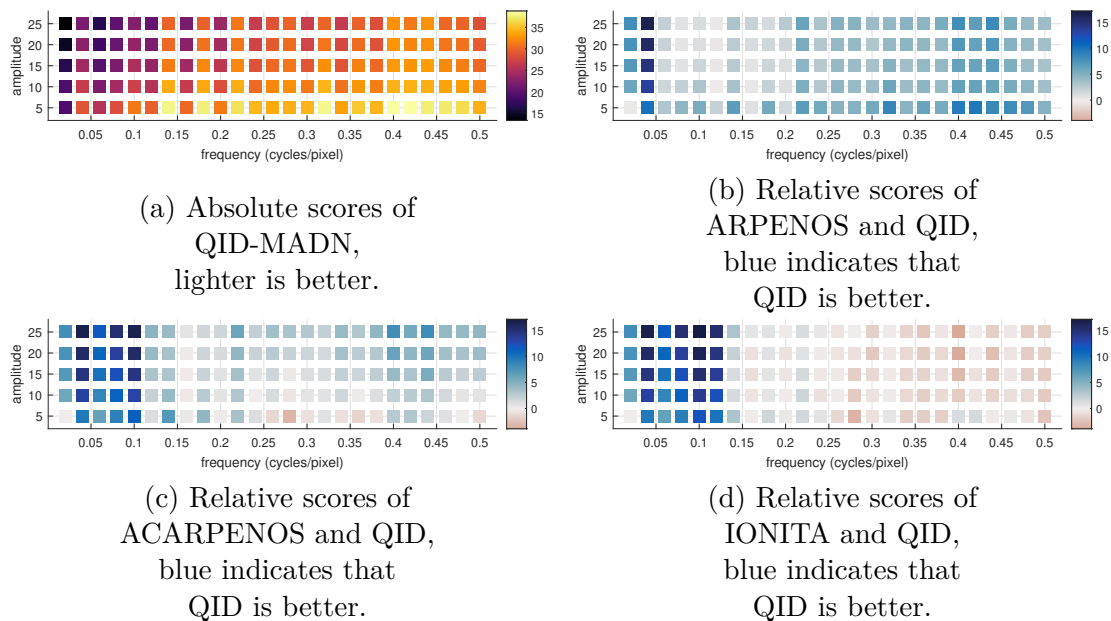


Figure 5-6: Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of SNR score.

Peak-signal-to-noise ratio (PSNR) is widely used for quality assessment of image compression algorithms. PSNR scores match almost exactly the SNR scores, but with an added offset. Their relative differences are almost negligible. No new comparative pattern is apparent for the PSNR metric, compared to SNR. Figures 5-7 and 5-8 show the identical patterns as discussed in Figures 5-5 and 5-6.

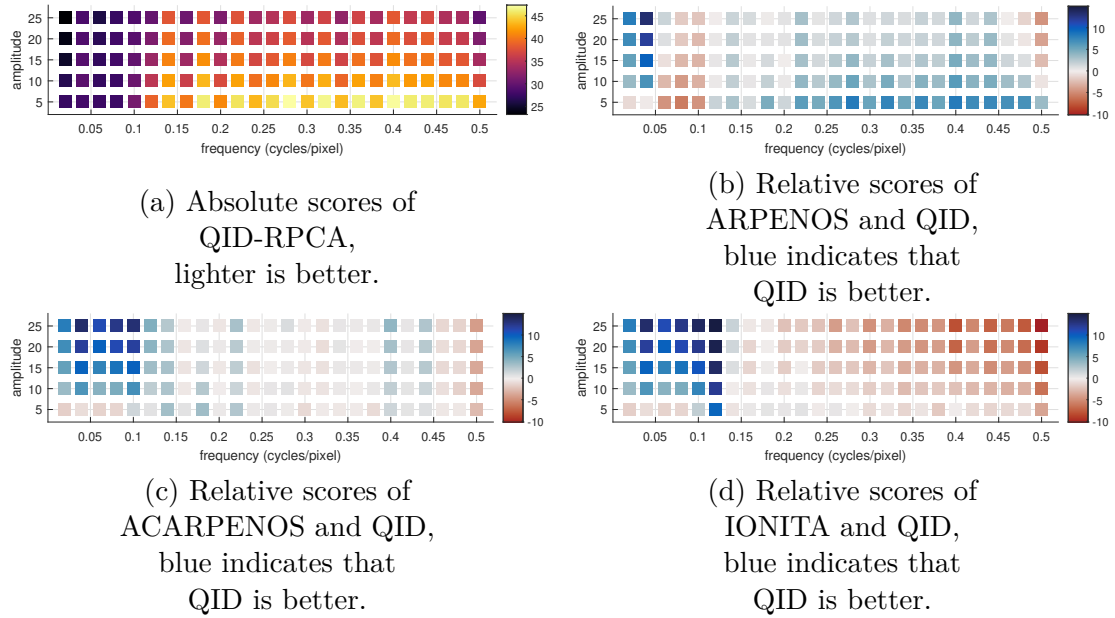


Figure 5-7: Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of PSNR score.

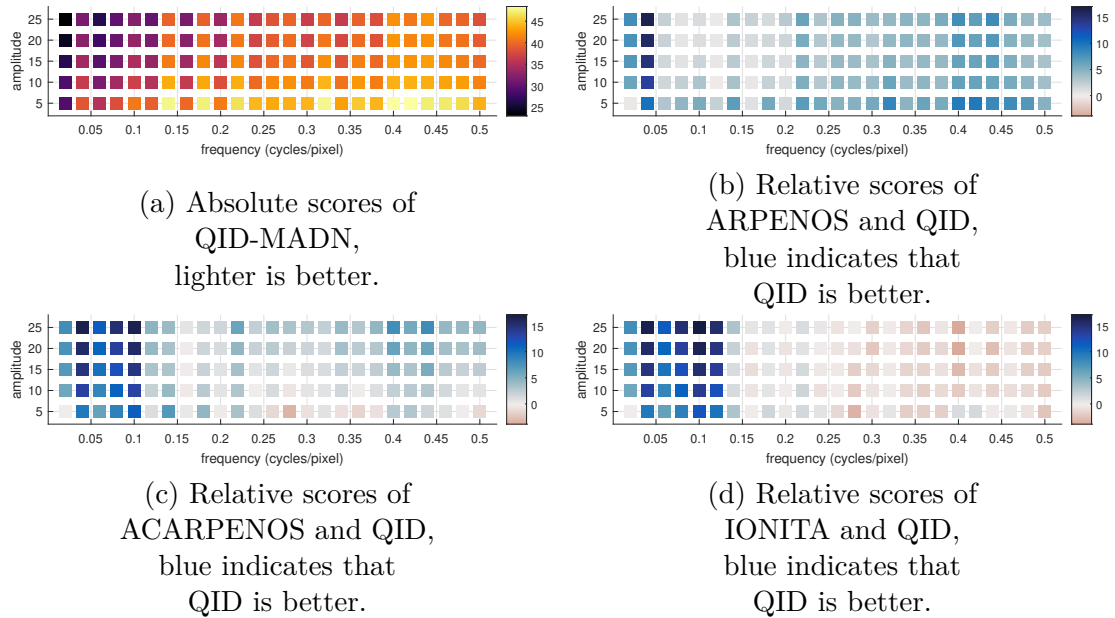


Figure 5-8: Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of PSNR score.

Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) is a no-reference image quality assessment metric that uses natural scene statistics (NSS). BRISQUE assumes that the undistorted image is a natural image, where the probability density function

of the normalized pixel intensities is close to a Gaussian distribution. Since the modified Shepp-Logan head phantom image is a synthetic, non-natural image and does not follow the aforementioned assumption, the BRISQUE results obtained in this case are not telling and should probably be disregarded. We are providing them here for completeness sake in Figures 5-9 and 5-10. However, when evaluating the real dataset, this no-reference metric might be a good indicator of image quality as long as its assumptions are not profoundly violated.

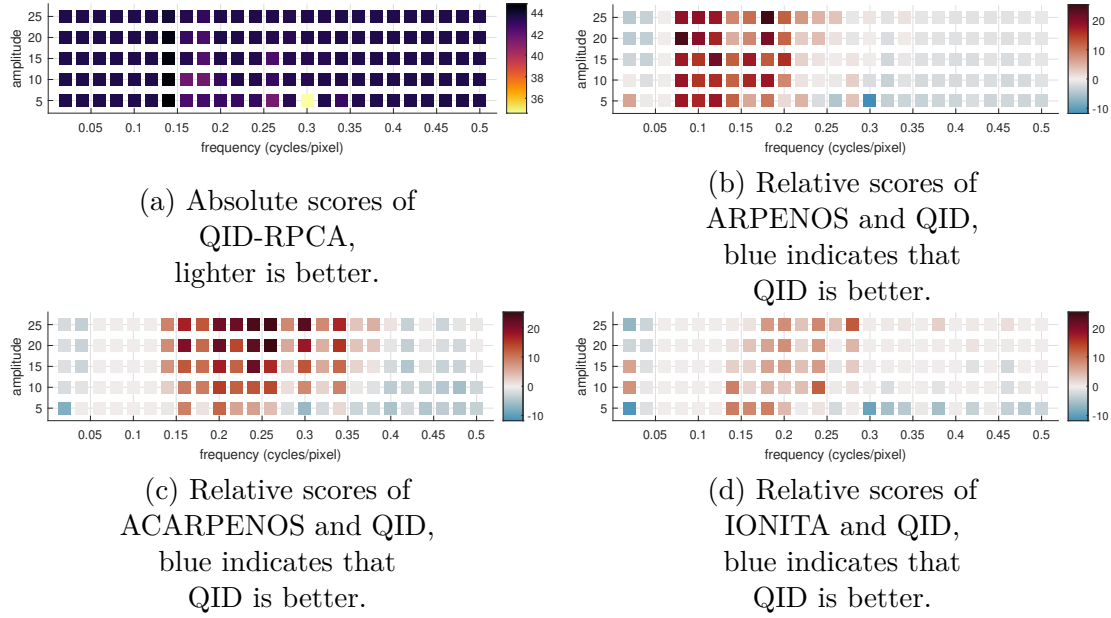


Figure 5-9: Differences between QID-RPCA, ARPENOS, ACARPENOS, and IONITA in terms of BRISQUE score.

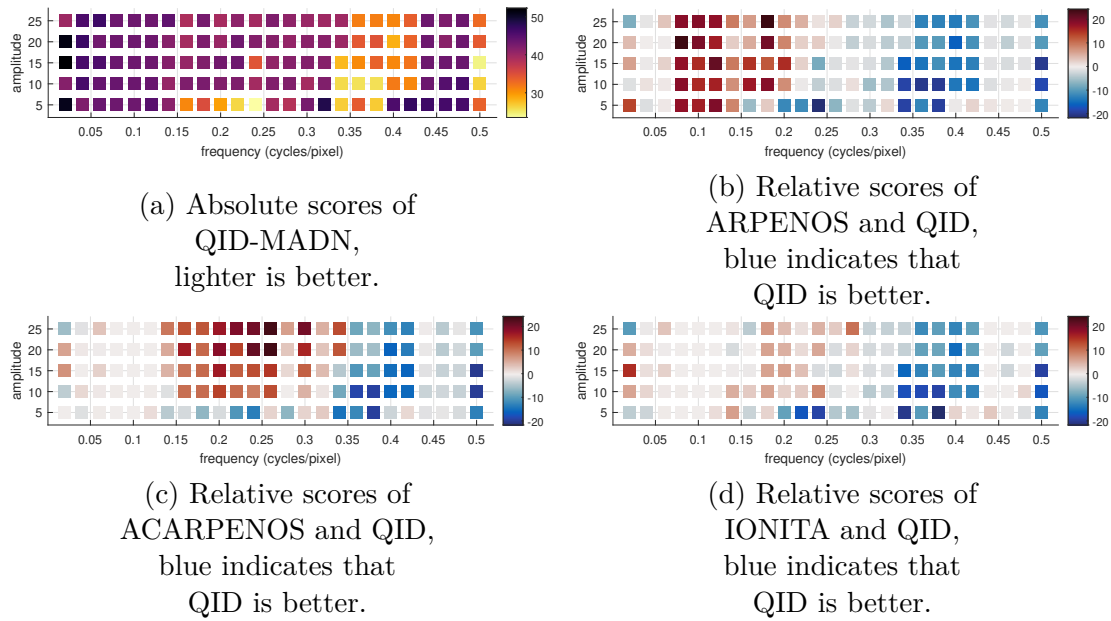


Figure 5-10: Differences between QID-MADN, ARPENOS, ACARPENOS, and IONITA in terms of BRISQUE score.

5-1-2. Image domain performance comparison of QID versus the state-of-the-art algorithms from literature

In this section, we provide image examples of the denoising performance and show examples of the visual effect of the largest score differences highlighted in Section 5-1-1. This is necessary, as the resulting images are expected to be inspected by humans. Consequently, denoising performance can be experienced as highly subjective, depending on the specific use-case. Evaluation of denoising performance such as "good", "bad", "better", and "worse" is done somewhat subjectively in this section with the main aim being to show the results so the reader can assess for themselves. Therefore, I want to note that my words here reflect my own perception and expectations towards high-quality denoising, as opposed to the previous section, where comparison was based on quantitative quality metrics.

Most of the following examples focus on the low-frequency performance of the algorithms, showing the breakdown points of specific algorithms. The difference in frequency of the introduced noise is of less importance than a specific algorithm's response to that noise frequency. One example will e.g. focus on the high-frequency performance differences between the algorithms. The last example e.g. highlights the difference in performance when the amplitude of the noise is small. This section focuses on the most prominent differences between the algorithms, these cases are not shown. The frequencies used are $[0.06, 0.08, 0.02, 0.5, 0.04]$ cycles/pixel.




Figure 5-11: Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.06 cycles/pixel, noise amplitude = 25 %.

As seen in Figure 5-11, ACARPENOS and IONITA are almost completely ineffective for such low-frequency corruptions. ARPENOS manages to mitigate the effects of the noise, but a considerable amount remains. The RPCA version of QID manages to eliminate the majority of the corruption. However, a strong artificial pattern, different from the original corruption, is introduced. This occurs because once an RPCA bin is identified to be an outlier, every coefficient within that bin is attenuated. For low-frequency bins, however, coefficient magnitudes change rapidly, and thus the expected value may not be a precise approximation for all the coefficients. Truncation to this imprecise approximation introduces new patterns. With regards to medical imaging, the acceptability of such a phenomena is application-specific. On the one hand, it could improve visual quality and help making the right diagnostic decision. On the other hand, introduced patterns may change the texture of the sample and mislead the expert, possibly resulting in misdiagnosis. The MADN version of QID, however, clearly delivers the highest quality denoising result in this comparison. The noise pattern is significantly reduced and there is no considerable artificial pattern introduced in the process. This behaviour can be explained by the more precise, individual filtering of coefficients, which only attenuates a select few coefficients within the disk.



Figure 5-12: Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.08 cycles/pixel, noise amplitude = 25 %.

Figure 5-12 shows a similar phenomenon as previously discussed in Figure 5-11. In this case, ARPENOS removes most of the corruption. However, it introduces an artificial

Gibbs pattern near the edge. The RPCA version of QID suffers from the same problem of introduced patterns for the low-frequency bins. The MADN version of QID removes a large majority of the noise, just like ARPENOS. However, QID-MADN still retains some residuals of the original corruption, even though introduction of artificial patterns is minimal.




Figure 5-13: Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.02 cycles/pixel, noise amplitude = 25 %.

As seen on Figure 5-13, none of the SOTA algorithms are active in this frequency range. The RPCA version of QID shows the same behaviour as discussed for Figures 5-11 and 5-12, with the main corruption pattern being eliminated, but a considerable amount of artificial noise being created, eroding the fine details and modifying the texture. QID-MADN, on the other hand, shows no signs of the same complications. The majority of the corruption is removed, while fine details are still clearly visible and the textures of uniform regions are largely unaltered compared to the ground truth image.




Figure 5-14: Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.14 cycles/pixel, noise amplitude = 25 %.

Figure 5-14 shows performance beyond the low-frequency limit of all the algorithms. All algorithms perform reasonably well in this range, with high quality removal of the corruption. IONITA and both versions of QID show residual patterns, while ARPENOS and ACARPENOS have minimal residuals. Introduction of artificial patterns is minimal, with some Gibbs phenomena visible where as the noise pattern is parallel with some edges of the image. However, textures and uniform areas are not affected significantly.




Figure 5-15: Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.5 cycles/pixel, noise amplitude = 25 %.

Figure 5-15 shows performance near the high-frequency limit of the algorithms, namely the Nyquist-frequency. Performance of the algorithms holds up, and there is little visible residual or artificial pattern introduced. For frequencies over this limit, denoising performance declines rapidly, more details are given in Section 5-1-4.




Figure 5-16: Comparison of the current state-of-the-art algorithms versus QID-RPCA and QID-MADN. Noise frequency = 0.04 cycles/pixel, noise amplitude = 5 %.

Most previous examples focused on high amplitude noise components. Figure 5-16, however, shows effectiveness for low amplitude patterns. Even though labeling of low-amplitude noise components is less straightforward, as peaks in the FFT do not stand out from neighboring coefficients as remarkably, detection and mitigation of such components is still done reasonably well by QID-MADN. The QID-RPCA version still suffers from loss of precision and introduction of artificial patterns in the low-frequency range.

5-1-3. Information loss through saturation of the dynamic range

Using the example of an 8-bit grayscale image, where each pixel can take a value between 0 and 255, values under or over those intensity limits cannot be represented. In cases where the added noise would result in intensity values outside of these boundaries, sensors often truncate these values to the closest representable value, 0 or 255 in this case. This breaks the superposition principle, the addition of noise is not linear anymore. This means that the corrupt image is not an additive sum of the clear, uncorrupted image and the noise pattern. This proves to be a very difficult denoising problem, as assumptions about the noise pattern are broken, and e.g. subtracting of the perfectly modelled quasiperiodic noise components still do not retrieve the original image.

In this case saturation results in a significant amount of information loss, which the al-

gorithms are unable to fully mend. Figures 5-17 and 5-18 portray denoising performance with information loss through saturation of the dynamic range. In the vertical axis, amplitude represents the amount of maximal truncation in % of the dynamic range. Even though it is clear, that both QID-MADN and QID-RPCA achieve low scores compared to a case where no information loss occurs, such as Figures 5-3 and 5-4, the results of the SOTA algorithms are comparably unsatisfactory. With the exception of the lower frequency range, where both QID-RPCA and QID-MADN have an advantage. Even though with scores so low, the difference hardly significant. Figure 5-19 shows the example where the difference is greatest between QID denoising and the SOTA. Even in this case, a considerable amount of residual patterns remain visible, with ARPENOS seemingly handling this type of information loss the best.

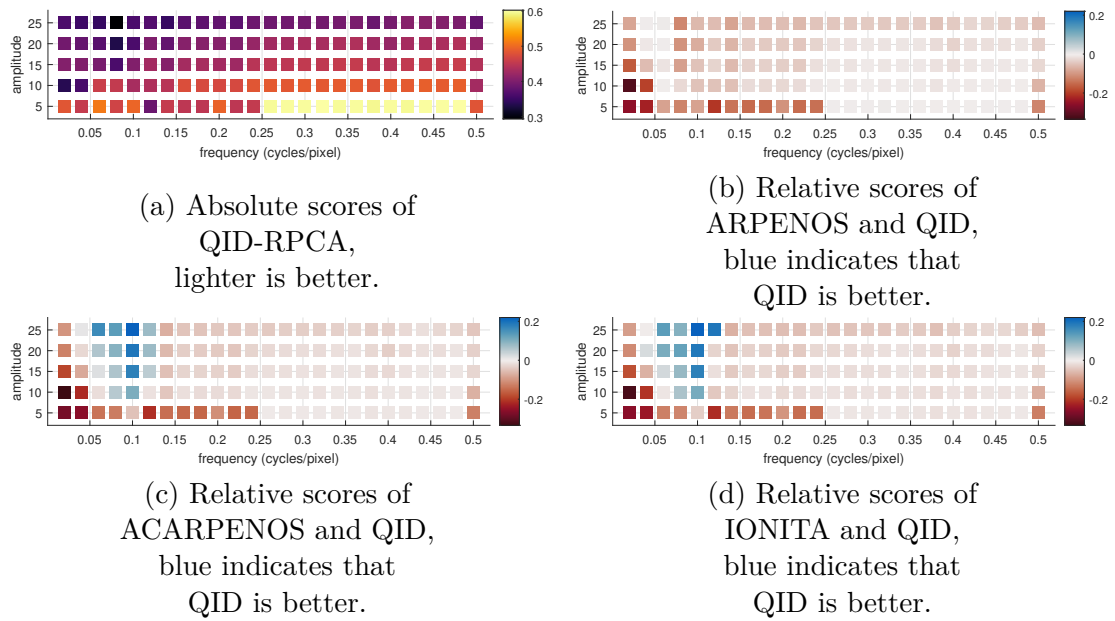


Figure 5-17: Effect of information loss through the saturation of the dynamic range of the image measured in terms of SSIM scores.

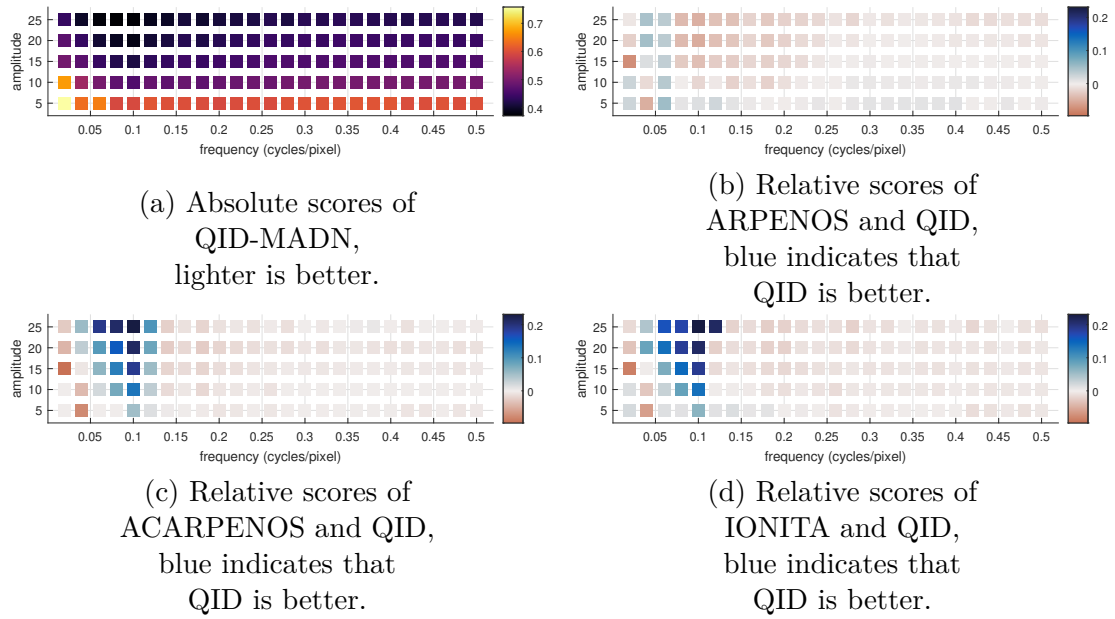


Figure 5-18: Effect of information loss through the saturation of the dynamic range of the image measured in terms of SSIM scores.



Figure 5-19: Comparison of the current State-of-the-art algorithms and QID under information loss. Noise frequency = 0.1 cycles/pixel, noise amplitude = 25 %.

5-1-4. Performance for denoising corruptions over the Nyquist-frequency

Each of the algorithms has an effective frequency bandwidth. Figures 5-20 and 5-21 show what occurs at and past the border of the bandwidth. Close to 0 cycles/pixel, all of these algorithms apply some form of protection of the zero-frequency component and its proximity. QID-MADN and QID-RPCA has it as low as 0.02 cycles/pixel, ARPENOS is set to 0.05 by default, while ACARPENOS and IONITA have their lower limit around 0.1 and 0.12 respectively. These limits heavily affect their effectiveness for low frequency corruptions. Most SOTA algorithms are very conservative regarding this limit, or would break down if this boundary would be lower, such as with ARPENOS where the linearity assumption would be heavily violated. However, QID tends to be flexible on this, while still offering effective protection of the zero frequency coefficient and it's neighborhood.

Looking at frequencies around and over 0.5 cycles/pixel QID-RPCA starts to get affected by the high-frequency limit, which is the frequency bandwidth upper limit given by the image's shorter spatial dimension. The frequency bandwidth is limited for QID-RPCA because mapping to a rectangular polar grid constrains the algorithm to the largest common radius in the 2D FFT, which is directly proportional to the shorter spatial dimension. This limit is 0.5 cycles/pixel, which, in the case of a 256 by 256 image, is equal to 128 cycles/width. When imposing noise components close to this border or above, some components tailing off the main spike may leak out of the detection bandwidth. Over this limit, QID-RPCA is inactive, thus strong residuals and purely noise related patterns may emerge.

QID-MADN and the SOTA algorithms are not constrained by this common radius, as these algorithms do not transform to the rectangular polar representation. These algorithms are still effective over this limit, as long as there are 2D FFT components representing the noise, which is the case if the noise components are not parallel with the horizontal or vertical axes, like in the synthetic dataset.

Figure 5-22 shows the performance difference between the QID and SOTA algorithms for the highest tested frequency. While the SOTA algorithms and QID-MADN perform well for high-frequency corruptions, QID-RPCA has a breakdown point at 0.5 cycles/pixel.

Overall, the performance of QID-RPCA declines for frequencies over its high-frequency limit given by the shorter spatial dimension of the image. QID-MADN and the SOTA algorithms do not suffer from the same limitation, and thus QID-MADN is equal in performance with SOTA algorithms for high-frequencies and performs significantly better for low-frequencies.

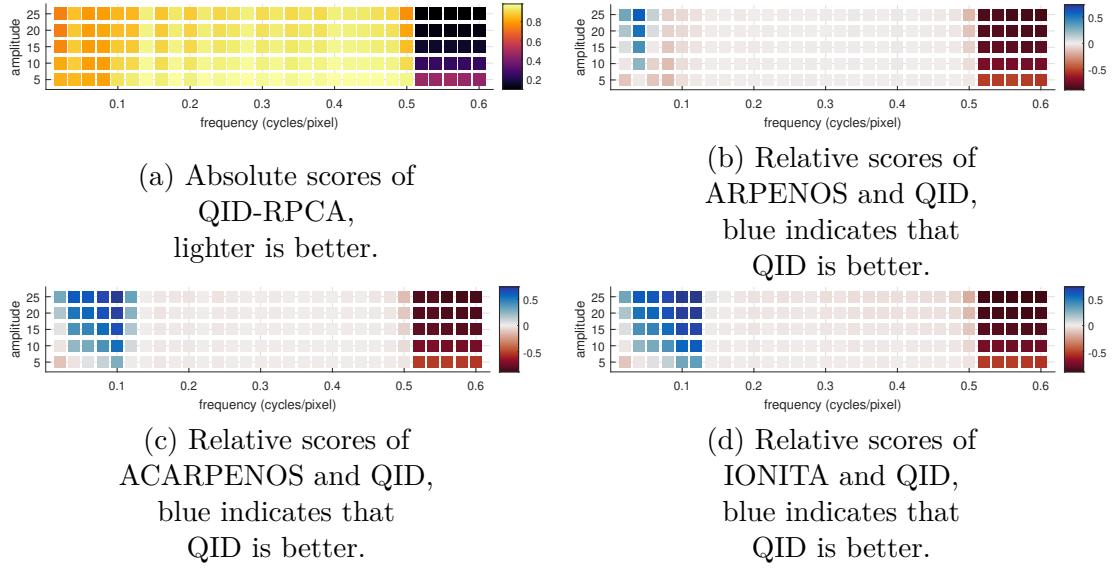


Figure 5-20: Effect of noise frequencies over and under cutoff frequency.

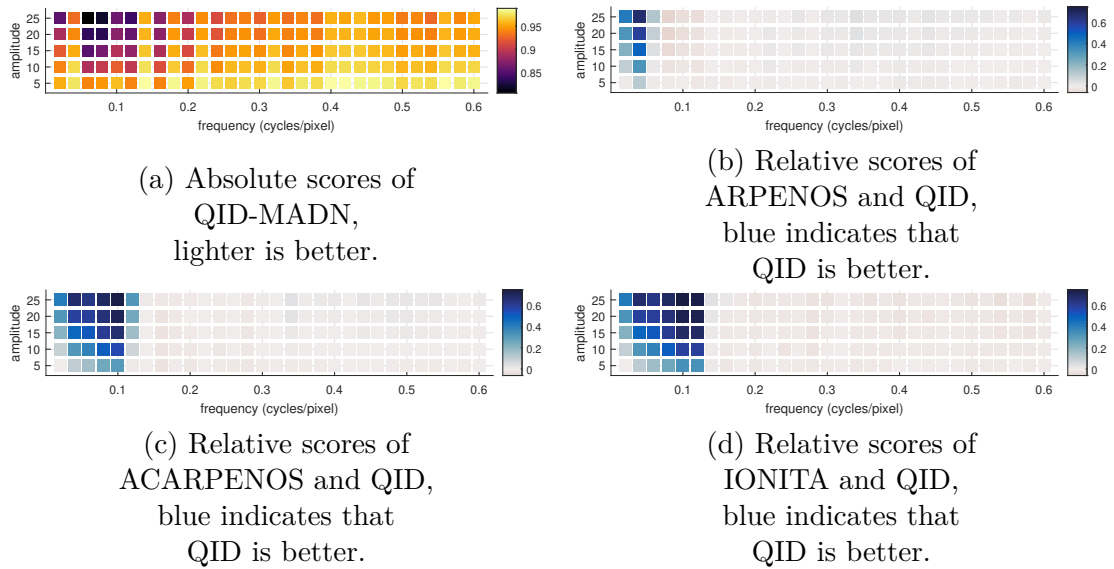


Figure 5-21: Effect of noise frequencies over and under cutoff frequency.




Figure 5-22: Comparison of the current State-of-the-art algorithms and QID for corruption over the Nyquist-frequency. Noise frequency = 0.6 cycles/pixel, noise amplitude = 25 %.

5-1-5. Runtime scaling

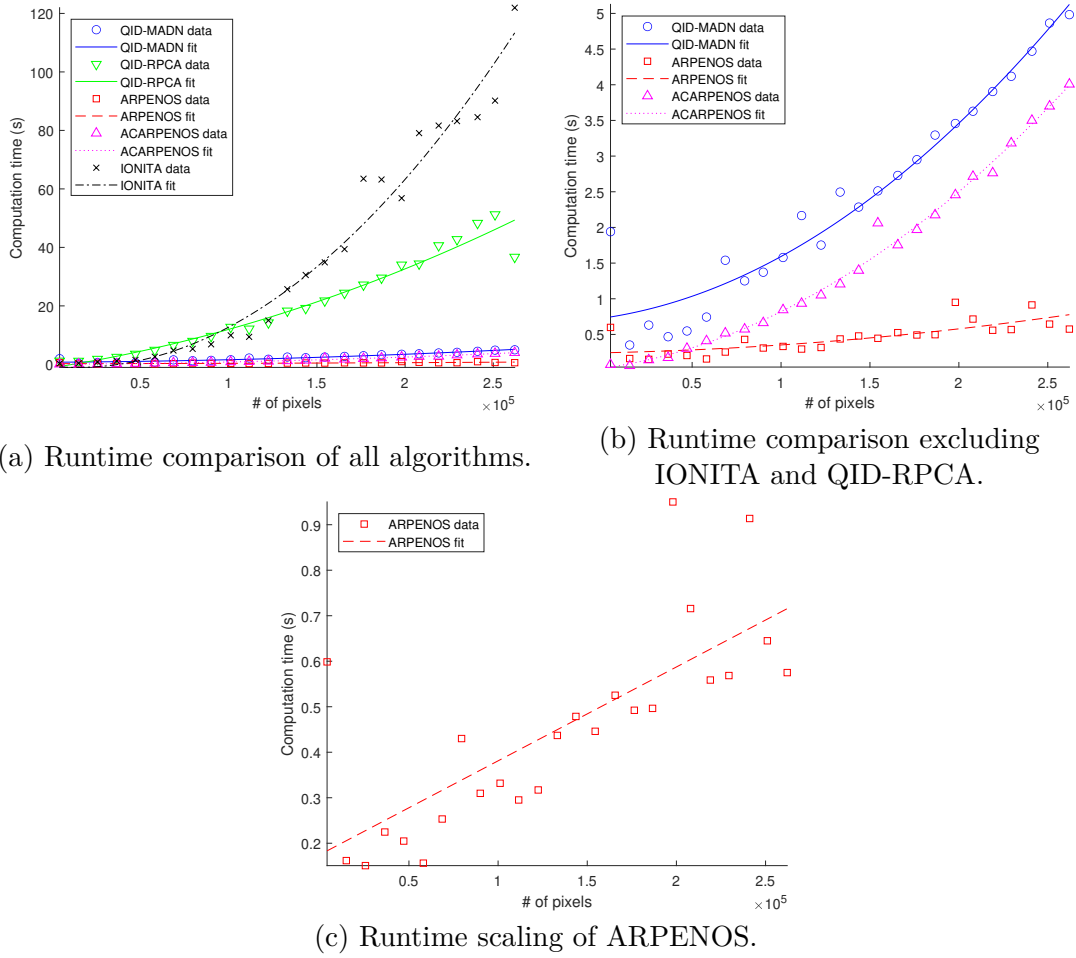


Figure 5-23: Plots showing computation times of each algorithm as a function of number of pixels.

The time scaling of IONITA is $O(n^2)$, as seen on the small scale test on Figure 5-23a, running on a two-core, 2.5 GHz processor. Extrapolating from this curve, computation time is over 9 hours for a 2048×2048 image, and approximately 30 minutes for a 1024×1024 image, which made further testing computationally impractical. QID-RPCA in the same scenario would take 2 hours for a 2048×2048 image, and approximately 9 minutes for a 1024×1024 image. This includes the optimization of the binning strategy, which effectively runs a large part of the calculations multiple times. ACARPENOS and QID-MADN are still best approximated by $O(n^2)$, however computation times are considerably lower (in seconds range), which makes them more applicable and practical for real-world use on large biomedical images. Given the simplicity and mostly vectorized computations of ARPENOS, it is able to maintain a low, close to $O(n)$, scaling, which can possibly be a very important factor if speed is of prime importance, while accuracy is secondary.

5-2. Processed dataset results and performance comparison

As described in Section 4-2, this dataset contains a labeled set of processed images that originate from single microscopy images of a transversal section of a mouse brain and a coronal section of a rat kidney. The processing includes different color spaces such as

RGB, HVS, and YCbCr, specific factors of decomposition and factorization methods such as PCA, NNMF, and DWT, range- and median filtering, entropy- and median filtering. The dataset also contains combinations of these, resulting in a vector of 905 distinct image bands with a varying level of noise and signal in each. A perceptually uniform color scale, *inferno*, has been used in order to promote uniformity and visibility of the images.

The exact preprocessing and modifications done to the image bands, hence the cause of the attenuation or amplification of a specific pattern is of secondary interest. This dataset, however, allows for testing and comparison of denoising performance on real-world images of biological origin while being corrupted by a varying range of artifacts. In this case, there is no objective, noiseless image, and as a result it is more difficult to give an unbiased comparison between the performance of the algorithms. However, a reference-free quality metric, BRISQUE, can be used, see in Section 2-5-2-4. Even though it will be biased towards a specific set of features, it provides a numerical basis for comparison. Subsequently, direct image comparisons are shown with the goal of highlighting specific behaviors and artifacts.

Tables 5-1 and 5-2 show the BRISQUE scores for each selected band of the mouse brain and rat kidney dataset. A lower score is better, but as it is demonstrated in Figures 5-24, 5-25, 5-26, and 5-27, a lower BRISQUE score does not necessarily mean that the denoising is of higher quality from the medical or human perspective. This discrepancy is partly because of the empty off-tissue buffer area around the biological sample and partly because clarity of biomedical images may not be modelled well by the natural scene statistics that BRISQUE measures.

It is clear from Tables 5-1 and 5-2, that the ACARPENOS results receives the same score as the raw image. ACARPENOS does not seem to be sensitive to corruptions of this kind. Element-by-element comparison to the raw image reveals that there is indeed no genuine denoising done by ACARPENOS. Based on the scores for the brain dataset, ARPENOS, IONITA, and QID-MADN perform best on average. For the kidney dataset, ARPENOS, QID-RPCA, and QID-MADN perform best on average. Not every band is shown and discussed in detail here since some bands show artifacts that are very similar to one another, both in the raw image and the denoised images.

Figure 5-24 shows the original and denoised versions of the 99th band of the processed mouse brain dataset. ACARPENOS is not sensitive to this corruption. ARPENOS, IONITA, and QID-RPCA remove a part of the noise pattern, but a significant residual is visible. ARPENOS also shows a blurring artifact, while QID-RPCA introduced a new artificial pattern. QID-MADN eliminates a large part of the original corruption and introduces no visible, artificial pattern.

Figure 5-25 shows the original and denoised versions of the 279th band of the processed mouse brain dataset. ACARPENOS is not sensitive to this corruption. ARPENOS, IONITA, and QID-RPCA remove part of the noise pattern, but a significant residual is still visible. QID-RPCA is the least effective out of these three algorithms. ARPENOS introduces a blurring artifact. QID-MADN eliminates a large part of the original corruption and introduces no visible, artificial pattern within the sample. Some patterns are introduced in the empty off-tissue buffer zone around the biological signal.

Figure 5-26 shows the original and denoised versions of the 99th band of the processed rat kidney dataset. ACARPENOS is not sensitive to this corruption. ARPENOS and QID-RPCA partially remove the noise, but residuals remain. QID-RPCA also introduced a new artificial pattern clearly visible in the buffer zone. IONITA and QID-MADN show high-quality denoising and introduce no visible, artificial patterns within the biological

sample.

Figure 5-27 shows the original and denoised versions of the 309th band of the processed rat kidney dataset. This band shows signs of intensity saturation in the low-end of the range. Saturation occurs where black spots are visible on the image. ACARPENOS is not sensitive to this corruption. ARPENOS and QID-RPCA partly mitigate the noise, but not completely. QID-RPCA seems to be of higher quality than ARPENOS, but introduced artificial patterns as well. IONITA and QID-MADN are of comparably high quality, with minor artificial patterns visible outside of the biological sample. Furthermore, the saturation artifact is partially ameliorated by IONITA and QID-MADN.

In summary, BRISQUE seems to be an insufficient metric for good quality assessment for the denoising of biomedical images. QID-MADN tends to perform better or at least equally well as the other algorithms. It was the only algorithm, that was steadily of high quality in all examples. However, QID-MADN only received the highest score in one out of the 14 cases. Other algorithms were inconsistent with their performance, sometimes provided a high-quality removal, but they also often introduced artificial patterns or contained a significant amount of residuals.

Table 5-1: BRISQUE scores of selected bands from the mouse brain dataset and images denoised by ARPENOS, ACARPENOS, IONITA, QID-RPCA, and QID-MADN.

Band # Algorithm	99	119	210	279	489	558	693	871
None, Raw image	37.95	37.96	41.48	38.84	43.66	44.83	44.54	43.38
ARPENOS	27.62	30.97	29.51	33.79	38.46	34.42	35.99	33.92
ACARPENOS	37.95	37.96	41.48	38.84	43.66	44.83	44.54	43.38
IONITA	20.52	24.35	21.10	29.67	40.22	28.13	17.09	26.45
QID-MADN	31.40	29.25	34.94	30.36	40.49	30.60	38.21	36.48
QID-RPCA	24.99	29.46	43.42	38.59	40.00	27.43	23.85	32.37

Table 5-2: BRISQUE scores of selected bands from the rat kidney dataset and images denoised by ARPENOS, ACARPENOS, IONITA, QID-RPCA, and QID-MADN.

Band # Algorithm	63	99	119	127	309	690
None, Raw image	45.59	42.18	42.67	44.92	42.86	41.90
ARPENOS	22.56	17.66	20.10	27.48	39.40	45.04
ACARPENOS	45.59	42.18	42.67	44.92	42.86	41.90
IONITA	42.82	37.26	19.73	31.21	33.19	39.97
QID-MADN	30.59	15.89	19.74	43.47	39.16	39.62
QID-RPCA	36.85	24.11	26.31	36.28	33.36	32.29

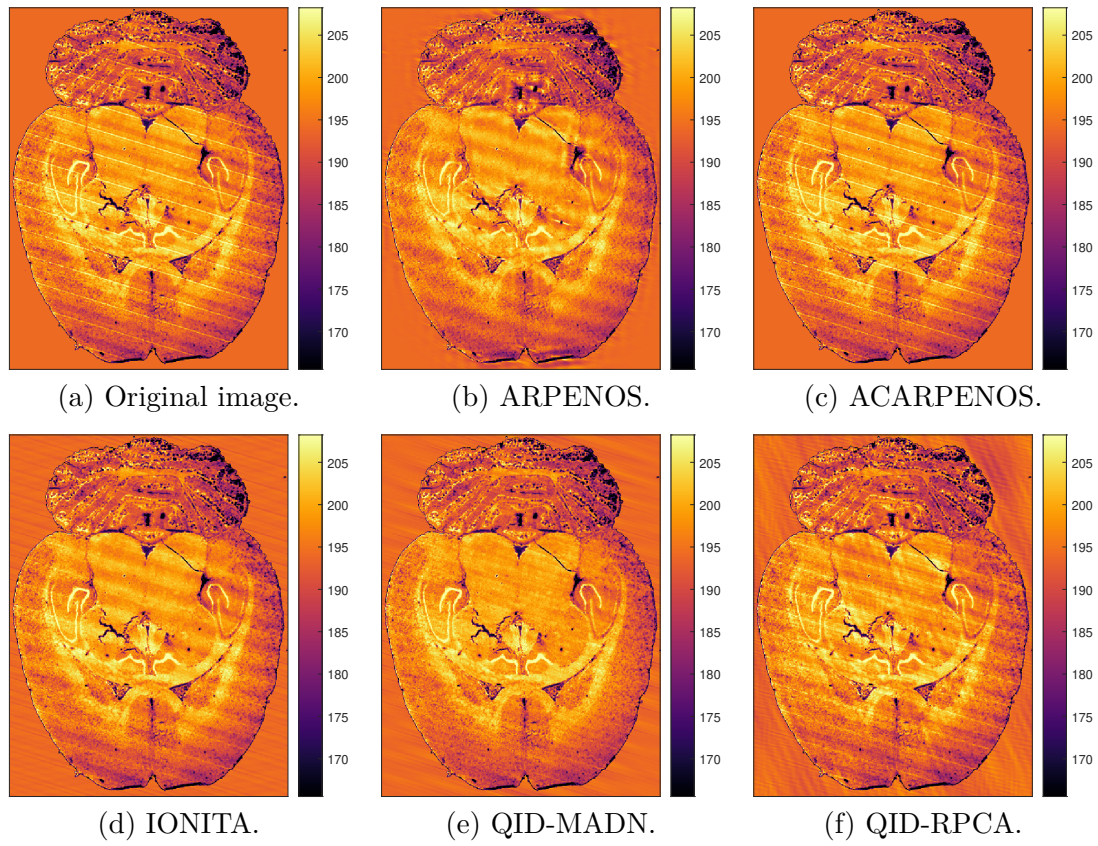


Figure 5-24: Noise removal of band 99 of the mouse brain dataset.

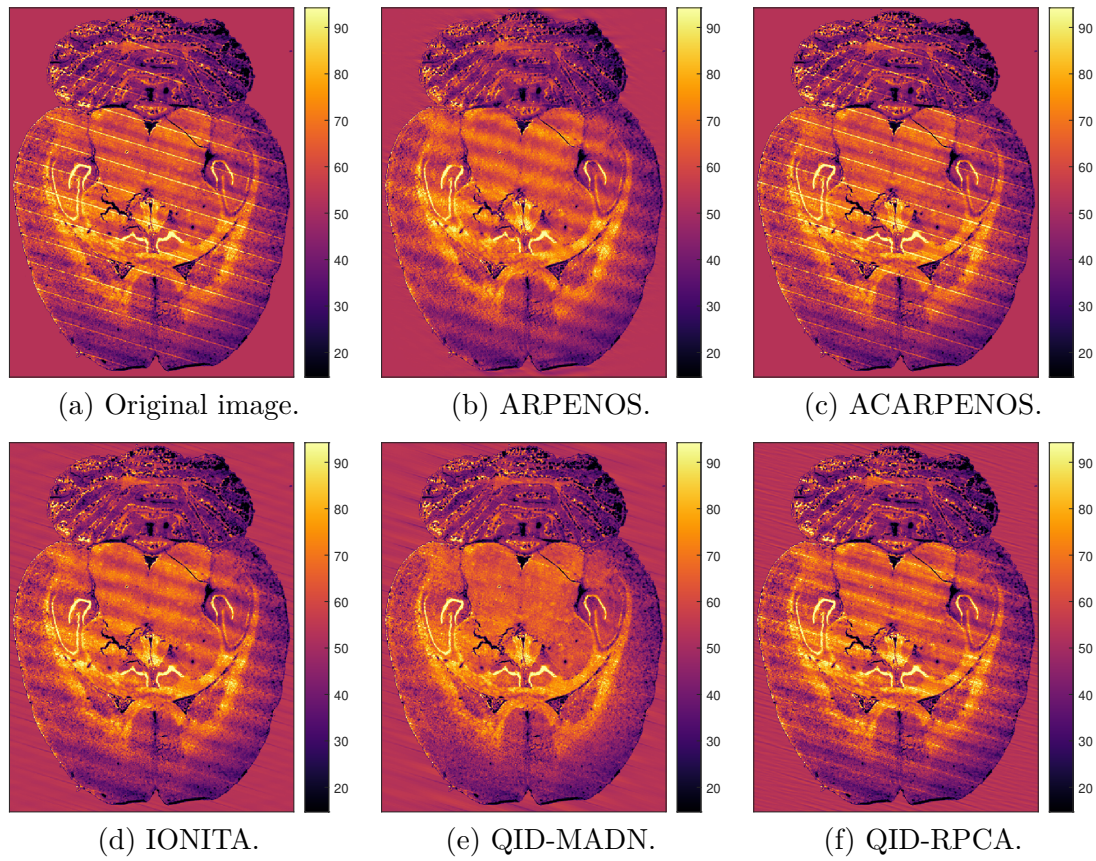


Figure 5-25: Noise removal of band 279 of the mouse brain dataset.

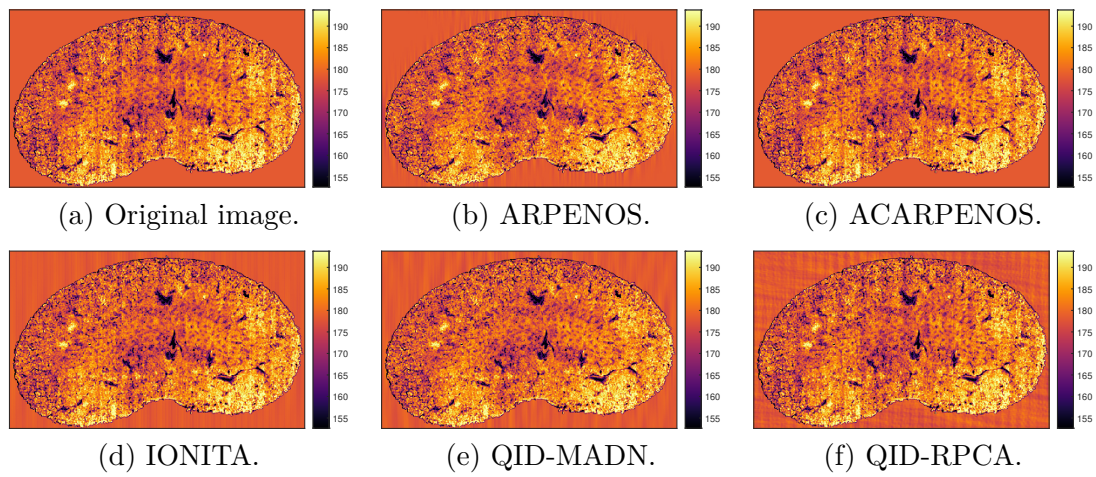


Figure 5-26: Noise removal of band 99 of the rat kidney dataset.

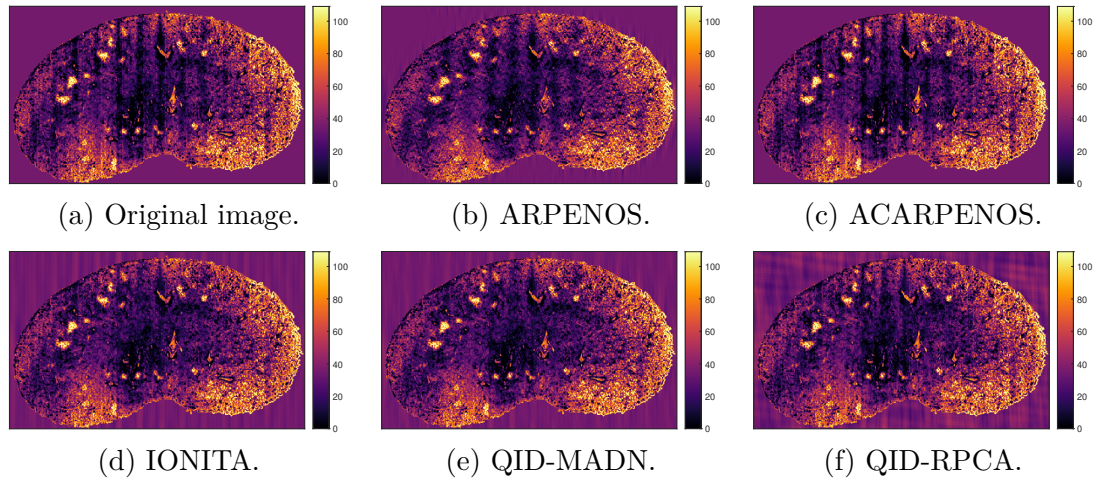


Figure 5-27: Noise removal of band 309 of the rat kidney dataset.

5-3. Real-world microscopy dataset results and performance comparison


Up until this point in the thesis, all experiments have been conducted on either synthetic images with controlled corruptions or on real-world microscopy images processed in such a way that certain patterns get amplified or attenuated. Even though, this provides valuable information on the strengths and weaknesses of the algorithms, it does not sufficiently showcase performance on the intended use-case: raw biomedical images.

This section explores the denoising performance on unaltered real-world microscopy images, exactly as they are acquired by the instrument. The algorithms are used without any preprocessing, knowledge of noise characteristics, or parameter tuning. For the SOTA algorithms, the parameters have been adjusted to the image dimensions in line with the recommendations in the literature. However, no further fine-tuning took place.


Details on the dataset and experimental setup can be found in Chapter 4-3. In short, a set of raw fluorescent microscopy images with varying degrees of quasiperiodic corruption are denoised by algorithms QID-MADN, QID-RPCA, ARPENOS, and ACARPENOS. Denoising using IONITA is omitted, as the runtime scaling of that specific algorithm made it impractical to use on high-resolution images, this is explained in more detail in Section 5-1-5. Each band in this dataset is a 27626×42238 image. Results of these denoising algorithms are compared and discussed in detail.

The dataset contains four separate bands, namely EGFP, DAPI, Cy5, and DsRed. Each of the original images are shown on Figures 5-28a, 5-29a, 5-30a, and 5-31a. Each pixel of these images is encoded as a 16-bit integer. However, the large majority of intensity values are contained in the lower end of the dynamic range, with very few and highly concentrated high-value pixels. As explained in Section 3-6, this reduces visibility to human eyes due to poor contrast. For this reason, the images have been manually saturated, as seen on Figures 5-28b, 5-29b, 5-30b, and 5-31b, in order to maximize visibility of the underlying biological information and corruption pattern. To stay true to the experimental setup though, this intensity saturation is only done for the visualizations, while the algorithms are still expected to work with the low intensity images and to handle poor contrast internally. In summary, the algorithms are used on the raw images, but saturated results are shown in the visualizations in this section. The colorbars on the right side of Figures 5-30b, 5-29b, 5-28b, and 5-31b indicate the exact saturation thresholds. These thresholds are

constant for a specific band throughout the analysis. A perceptually uniform color scale, *inferno*, has been used in order to promote uniformity and visibility of the images.




(a) Raw microscopy image.




(b) Microscopy image saturated for increased visibility.

Figure 5-28: EGFP band of the human kidney fluorescent microscopy dataset.




(a) Raw microscopy image.




(b) Microscopy image saturated for increased visibility.

Figure 5-29: DAPI band of the human kidney fluorescent microscopy dataset.




(a) Raw microscopy image.




(b) Microscopy image saturated for increased visibility.

Figure 5-30: Cy5 band of the human kidney fluorescent microscopy dataset.



(a) Raw microscopy image.




(b) Microscopy image saturated for increased visibility.


Figure 5-31: DsRed band of the human kidney fluorescent microscopy dataset.

As an example, Figure 5-32 shows the denoised images for the EGFP band. EGFP and other bands are discussed in more detail in the following sections. At first glance, there is no visible difference between the original image, the ARPENOS, and ACARPENOS denoised image. QID-MADN and QID-RPCA, however, visibly reduced the effects of viewport stitching. Element-by-element comparison of the original, ARPENOS, and


ACARPENOS denoised images confirm that they are indeed identical for each separate band. This is likely caused by the frequency of the stitching pattern being under the detection threshold of these algorithms and within the protection area around the zero-frequency. Further analysis of this case study only focuses on comparisons between the original and the QID denoised images, in order to reduce the number of redundant images for clarity. However, the same remarks apply for the original, ARPENOS, and ACARPENOS.




(a) Saturated microscopy image.




(b) Denoising by ARPENOS.



(c) Denoising by ACARPENOS.



(d) Denoising by QID-MADN.



(f) Denoising by QID-RPCA.

Figure 5-32: Noise removal of EGFP band of the human kidney fluorescent microscopy dataset.

5-3-1. EGFP band, in-depth analysis

Figure 5-33 shows the denoised images. There is no visible difference between the original image, the ARPENOS, and ACARPENOS denoised image. QID-MADN and QID-RPCA, however, visibly reduced the effects of viewport stitching. QID-RPCA also introduced an unsatisfactory amount of artificial patterns similar to speckle noise.


The stitching pattern is most prevalent in the central portion of the image, even though, it

is visible throughout the image, both within the tissue sample and in the empty recorded off-tissue area around the sample referred to as the buffer zone. The only exception is the mask area, where no physical measurements were acquired, which was added artificially to make the image rectangular. The QID-MADN version of the algorithm mitigates the viewport stitching pattern significantly, minor artificial patterns are introduced, such as the 'halo' effect in the empty buffer zone. Furthermore, a strong contrast around the edge of the sample appears. It is most recognizable where the edge of the sample and the horizontal or vertical pattern of the viewport pattern are parallel. QID-RPCA displays the same minor patterns, yet, a strong, speckle-like, dotted pattern appears as well, which is comparable in strength to the original viewport-stitching. This pattern reduces the denoising performance significantly. QID-MADN, on the other hand, presents a generally high-quality denoised result, where the reduction in noise is significant and the introduced patterns are minor.


Figure 5-34 shows a zoomed-in section of the central portion of the EGFP band of the image. This portion contains multiple, visually distinguishable tissue structures and biological regions. For QID-RPCA, the speckle pattern is prevalent and distinction of biological structures is problematic, as the grain size of the introduced pattern is similar to patterns of biological origin. After denoising with QID-MADN, the tissue structures in these regions look more uniform, while still being easily distinguishable. No visible loss of detail is apparent and it is possibly easier to recognize irregular behaviour within these uniform regions. A minor concern is the introduced low-intensity spots, specifically in the vicinity of the edges of the sample where the average intensity is high.

Figure 5-35 shows the bottom-right corner of the image. Both algorithms show a significant change in contrast around the edge of the sample, where the noise pattern is parallel to it. The QID-RPCA introduced pattern is similar in size to the biological pattern, which makes distinction of artificial and biological patterns problematic. QID-MADN portrays high-quality denoising. However, there is a noticeable intensity difference between the uniform tissue patches between the small, circular high-intensity areas. This is the result of the minimally invasive denoising strategy. Even though it may be visible largely due to the aggressive saturation, it is important to be cautious of the medical implications of such patterns. Around the horizontal edge of the sample a strong presence of what seems to be Gibbs-phenomena, which are introduced by QID-MADN.


Figure 5-36 shows a zoomed-in section of the center of the power spectrum on a logarithmic scale, as well as the outlier mask and noise component attenuation done by QID-MADN. The power coefficients themselves are not trivial to interpret. However, since the viewport stitching pattern is relatively low frequency and has well defined directions along the horizontal and vertical axes, the coefficients responsible are effectively identified and attenuated by the algorithm. The diagonal lines likely correspond to diagonal edges of the sample, while the individual peaks correspond to the a more overarching pattern, such as the viewport-stitching. Note the individual peaks in the outlier map and the corresponding spectral component along the horizontal and vertical axis.



(a) Saturated microscopy image.




(b) Denoising by QID-MADN.




(c) Denoising by QID-RPCA.


Figure 5-33: Noise removal of EGFP band of the human kidney fluorescent microscopy dataset.



(a) Saturated microscopy image.




(b) Denoising by QID-MADN.




(c) Denoising by QID-RPCA.


Figure 5-34: Zoomed-in portion of the central part of the EGFP band of the human kidney fluorescent microscopy dataset.



(a) Saturated microscopy image.



(b) Denoising by QID-MADN.



(c) Denoising by QID-RPCA.

Figure 5-35: Zoomed-in portion of the right corner of the EGFP band of the human kidney fluorescent microscopy dataset.




Figure 5-36: Zoomed-in central portion of the EGFP band of the logarithm of the raw power spectrum, the outlier map, and the logarithm of the filtered power spectrum.


5-3-2. DAPI band, in-depth analysis

As discussed before, there is no visible difference between the original image, the ARPENOS, and ACARPENOS denoised image, and so this comparison will be omitted here to focus on QID. Furthermore, QID-RPCA shows the same unsatisfactory amount of artificial patterns as discussed in Section 5-3-1. This is due to some false positive labeling of some higher-frequency bins, as the outlier detection is not sufficiently robust. Since this behavior is the same as discussed before, images of this algorithms are omitted here as well.


The comparison between the raw image and QID-MADN is shown on Figure 5-37. The denoising performance of QID-MADN reduces significantly for this example. The quasiperiodic viewport-stitching pattern is partly eliminated, but significant artificial patterns are introduced as well. High-intensity regions are lowered, while low intensity regions are raised in value. This behaviour may be explained by the power spectrum in Figure 5-38, where distinct directions and peaks are hard to identify even manually. The artificial patterns are likely caused by the components just outside of the zero-frequency protection area. It seems that these components do not correspond to the horizontal and vertical peaks, seen in the previous example, which causes false positive outliers and added artificial

patterns. This bin covers the largest radius and is expected to contain the largest range of genuine, noise-free power spectrum values. A more granular radial binning in this specific region may alleviate some false positives. Furthermore, the viewport stitching on the non-saturated image has very low amplitude, as seen on Figure 5-29a. This seems to be just under the detection threshold and may only be visible because of the applied saturation, as opposed to the EGFP band, where the pattern is more clearly distinguishable.

The DsRed band illustrates the same phenomena, and the same explanation can be given for its sub-optimal performance. Therefore, this example is not discussed here in detail.



(a) Saturated microscopy image.



(b) Denoising by QID-MADN.

Figure 5-37: Noise removal of DAPI band of the human kidney fluorescent microscopy dataset.





Figure 5-38: Zoomed-in central portion of the DAPI band of the logarithm of the raw power spectrum, the outlier map, and the logarithm of the filtered power spectrum.

5-3-3. Cy5 band, in-depth analysis


As discussed before, there is no visible difference between the original image, the ARPENOS, and ACARPENOS denoised image, and so this comparison will be omitted here. Furthermore, QID-RPCA shows the same unsatisfactory amount of artificial patterns as discussed in Section 5-3-1. This is due to some false positive labeling of some higher-frequency bins. Since this behavior is the same as discussed before, images of this algorithm are also omitted.

In Figure 5-41, QID-MADN portrays high-quality denoising. However, there is a noticeable intensity increase for the uniform tissue regions between the high-intensity areas, most visible in the bottom left of the image. Even though this may be visible largely due to the aggressive saturation of these visualizations, it may also be explained by the outliers closest to the zero-frequency, visible on Figure 5-42b. The power spectrum visible on Figure 5-42a shows easily distinguishable peaks and directions of high power, which are identified and filtered by QID-MADN.

Interestingly, the central region shown on Figure 5-40 is affected considerably less by this pattern and is generally denoised well.




(a) Saturated microscopy image.




(b) Denoising by QID-MADN.

Figure 5-39: Noise removal of Cy5 band of the human kidney fluorescent microscopy dataset.




(a) Saturated microscopy image.




(b) Denoising by QID-MADN.

Figure 5-40: Zoomed in portion of the central part of the Cy5 band of the human kidney fluorescent microscopy dataset.



(a) Saturated microscopy image.



(b) Denoising by QID-MADN.

Figure 5-41: Zoomed in portion of the right corner of the Cy5 band of the human kidney fluorescent microscopy dataset.




Figure 5-42: Zoomed-in central portion of the Cy5 band of the logarithm of the raw power spectrum, the outlier map, and the logarithm of the filtered power spectrum.

5-3-4. Summary

Overall, this fluorescent microscopy case study demonstrates that this viewport-stitching pattern is under the detection threshold of algorithms ARPENOS and ACARPENOS, and that both QID algorithms manage to mitigate it. Even though the initial corruption is removed, denoising done by QID-RPCA is dissatisfactory, with the speckle pattern it tends to introduce being considerable. Likely, the pattern appears because the binning strategy is least precise for low-frequencies. Attenuating every component in the bins that contain the most important noise components introduces substantial artificial patterns. Furthermore, the non-robust filtering approach gives rise to higher-frequency false positives. However, denoising done by QID-MADN is of high-quality. The initial corruption is removed and apart from minor artifacts, minimal change has been introduced into the biological information. In cases where the corruption pattern has low amplitude and power, some patterns may fall under the detection threshold. Furthermore, since precision is lowest right outside the border of the zero-frequency protection area, false positive outliers may be identified, which can introduce relatively minor artificial patterns.

Section 5-1 shows that the detection limits of both QID algorithms are considerably more

generous compared to the SOTA. This is especially true for low-frequency variation. Performance of all the algorithms are comparable for mid-frequency corruptions. For high frequencies, around and over 0.5 cycles/pixel, QID-MADN and SOTA are comparable in scores, while QID-RPCA has a breakdown point exactly at 0.5 cycles/pixel. Information loss through saturation affects all algorithms negatively. QID-MADN performs consistently at the same rate, or better than the next best performing algorithm even in this case.

Section 5-2 shows examples of processed microscopy images of a mouse brain and rat kidney. It shows that the BRISQUE no-reference metric may be insufficient for biological images which have a relatively large mask without biological information. ARPENOS, ACARPENOS, IONITA, and QID-RPCA mitigated the noise, but often had considerable residuals, especially in the lower-frequencies. QID-MADN consistently outperformed or had an equal denoising performance as the next best algorithm.

Section 5-3 illustrates the expected behaviour on raw, high-resolution microscopy images. The images also have low contrast. ARPENOS, ACARPENOS, and IONITA are inactive for the type of quasiperiodic corruption that is present here and do not modify any of the bands. QID-RPCA and QID-MADN remove a large majority of the viewport-stitching pattern. Furthermore, QID-RPCA introduces a significant speckle-like pattern, which is dissatisfactory. QID-MADN, however, delivers the highest quality denoising in the comparison, with minimal introduction of artificial patterns.

In conclusion, the current SOTA algorithms and QID-RPCA are consistently outperformed by QID-MADN. Denoising of quasiperiodic patterns in biomedical images in a data-driven manner, without the need for manual parameter tuning or preliminary knowledge of the noise parameters is possible and viable, and we believe QID-MADN can be an important factor in bringing this capability into the life sciences practice.

Chapter 6

Conclusions and future work

6-1. Conclusions

In this thesis, we developed a fully data-driven, quasiperiodic noise filtering framework for biomedical images, QID. Quasiperiodic corruptions are repeating patterns that are of non-biological origin, such as instrument or acquisition method-induced noise. Some of these quasiperiodic patterns include e.g. viewport stitching artifacts, slice-to-slice interference, aliasing, and Gibbs-phenomena. Separating the biological signal from noise is of significant importance (e.g. for faster decision-making, reduced chance of misdiagnosis), as medical decisions may be based on the results of biomedical imaging.

Previously developed algorithms, such as ARPENOS, ACARPENOS, and IONITA have shortcomings such as manual tuning of multiple parameters, ineffectivity for low-frequency corruptions, strong assumptions about the underlying noise-free signal, and high computation times especially for high-quality, high-resolution images. Our aim was to build the novel framework on the strengths of these algorithms, and at the same time to expand it in such a way that the aforementioned shortcomings are somewhat alleviated.

Two different algorithms were implemented within the framework, QID-RPCA and QID-MADN. One algorithm is based on robust principal component analysis (QID-RPCA) and the other uses the normalized median of absolute differences (QID-MADN) to determine outlier coefficients. The methods chapter details a number of steps, including the preprocessing, noise component labeling and attenuation, and image restoration. This chapter also includes the reasoning and motivation behind the design choices that facilitate the unsupervised, data-driven denoising algorithms. Both algorithms operate in the Fourier (or frequency) domain, and the steps combine e.g. the use of histogram equalization for radial binning, an automated, sparsity-based approach to choosing the optimal aggregation level, and noise component attenuation based on radial patterns in the power spectrum.

The case study chapter introduces three datasets and outlines the set of experiments performed on them. A synthetic dataset based on the Shepp-Logan phantom is corrupted by a wide range of periodic corruptions, both in terms of frequency and amplitude. Comparing the denoising performance of the state-of-the-art algorithms and QID highlights the strengths and weaknesses of the algorithms as well as the limits of their applications. Then, the methodology is demonstrated on two real-world datasets, one based on a microscopy image of a transversal section of a mouse brain and the other one is a microscopy

image of a coronal section of a rat kidney. Finally, the algorithms are applied on a set of real-world, high-resolution fluorescent microscopy images of a human kidney. This dataset illustrates the performance on raw, completely unaltered biomedical images.

In conclusion, the goal of this thesis was to improve on and, where possible to surpass current state-of-the-art quasiperiodic denoising algorithms. This was done through the development of an automated, entirely data-driven framework, QID, that is able to reliably identify, quantify, and eliminate quasiperiodic, non-biological patterns within the images while retaining as much biological information as possible. Two novel algorithms were implemented, QID-RPCA and QID-MADN. Results indicate that the novel algorithms have higher denoising performance than previous approaches in the literature with notable improvements achieved for low-frequency corruptions. This is achieved without the need for manual parameter tuning, strong assumptions about the noise-free signal, or computation times that would obstruct wide-scale applicability. One of the algorithms, QID-MADN, consistently outperformed the other algorithms, by a considerable margin, while the other version, QID-RPCA, is susceptible to false positives and the introduction of artificial patterns with its current implementation. Overall, the initial goals have been achieved and a promising framework has been implemented for removing quasiperiodic corruptions from biomedical images in an unsupervised fashion.

6-2. Future work

This section list a few vectors for future research. Implementing these ideas can potentially further improve the unsupervised denoising performance on quasiperiodic corruptions.

- Automation of padding and apodization can improve denoising quality and reduce residual patterns. However, careful modeling would be required in order to guarantee that no artificial patterns are introduced through the process.
- Changing the QID-RPCA approach to use a more robust outlier labeling approach, could decrease its rate of false positives, and thus reduce the amount of introduced patterns. The current approach does not allow median based labeling as the exact fit property is often satisfied.
- The phase component of the Fourier representation of an image encodes the location of image features, such as lines and edges. Expanding the analysis to the phase components, rather than focusing solely on the magnitude components as we've done here, can potentially aid keeping important image features intact throughout the denoising process.
- Expanding the preprocessing to incorporate contrast stretching or histogram equalization to handle low-contrast images. Outliers in these images are more likely to fall under the detection limit of the algorithm as the power of the components is often indistinguishable from non-noise components within in the same radius.
- Finding the optimal binning strategy for QID-RPCA currently requires computationally costly steps to reach the sparse component. Finding a good predictor of the quality of the binning strategy and its sparsity with a reduced computational complexity may help the practical use of the algorithm significantly.
- For QID-MADN, a more granular binning near the central components may be beneficial. This may increase precision for high-resolution images with low-frequency corruptions. Furthermore, it would not increase computation times significantly as it only affects the low-frequency region. Since median-based filtering approaches are

less sensitive to different sample sizes, the reliability of the outlier labeling is not expected to be affected significantly.

Bibliography

- [1] J. B. J. Fourier, Theorie analytique de la chaleur. Paris: F. Didot, 1822, 639 pp.
- [2] A. G. Bell, "Upon the production and reproduction of sound by light", Journal of the Society of Telegraph Engineers, vol. 9, no. 34, pp. 404–426, 1880. DOI: [10.1049/jste-1.1880.0046](https://doi.org/10.1049/jste-1.1880.0046).
- [3] Kendall M. G., An Introduction To The Theory Of Statistics. 1911. [Online]. Available: <http://archive.org/details/in.ernet.dli.2015.233345>.
- [4] H. Dalton, "The measurement of the inequality of incomes", The Economic Journal, vol. 30, no. 119, pp. 348–361, Sep. 1, 1920, Publisher: Oxford Academic, ISSN: 0013-0133. DOI: [10.2307/2223525](https://doi.org/10.2307/2223525). [Online]. Available: <https://academic.oup.com/ej/article/30/119/348/5282305>.
- [5] S. H. Gage, "Modern dark-field microscopy and the history of its development", Transactions of the American Microscopical Society, vol. 39, no. 2, p. 95, Apr. 1920, ISSN: 00030023. DOI: [10.2307/3221838](https://doi.org/10.2307/3221838). [Online]. Available: <https://www.jstor.org/stable/3221838?origin=crossref>.
- [6] C. Gini, "Measurement of inequality of incomes", The Economic Journal, vol. 31, no. 121, pp. 124–126, 1921, Publisher: [Royal Economic Society, Wiley], ISSN: 0013-0133. DOI: [10.2307/2223319](https://doi.org/10.2307/2223319). [Online]. Available: <https://www.jstor.org/stable/2223319>.
- [7] M. von Ardenne, "Das Elektronen-Rastermikroskop", Zeitschrift für Physik, vol. 109, no. 9, pp. 553–572, Sep. 1, 1938, ISSN: 0044-3328. DOI: [10.1007/BF01341584](https://doi.org/10.1007/BF01341584). [Online]. Available: <https://doi.org/10.1007/BF01341584>.
- [8] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex fourier series", Mathematics of Computation, vol. 19, no. 90, pp. 297–301, 1965, ISSN: 0025-5718. DOI: [10.2307/2003354](https://doi.org/10.2307/2003354). [Online]. Available: <https://www.jstor.org/stable/2003354>.
- [9] L. Amerio and G. Prouse, Almost-Periodic Functions and Functional Equations, ser. The university series in higher mathematics. New York: Springer-Verlag, 1971, ISBN: 978-1-4757-1256-8. DOI: [10.1007/978-1-4757-1256-4](https://doi.org/10.1007/978-1-4757-1256-4). [Online]. Available: <https://www.springer.com/gp/book/9781475712568> (visited on 07/04/2020).
- [10] L. A. Shepp and B. F. Logan, "The fourier reconstruction of a head section", IEEE Transactions on Nuclear Science, vol. 21, no. 3, pp. 21–43, Jun. 1974, Conference Name: IEEE Transactions on Nuclear Science, ISSN: 1558-1578. DOI: [10.1109/TNS.1974.6499235](https://doi.org/10.1109/TNS.1974.6499235).
- [11] G. N. Hounsfield, "Method of and apparatus for examining a body by radiation such as x or gamma radiation", Originating Research Org. not identified, US 3919552, Nov. 11, 1975. [Online]. Available: <https://www.osti.gov/biblio/4102886->

- [method-apparatus-examining-body-radiation-gamma-radiation](#) (visited on 07/08/2020).
- [12] A. Jerri, "The shannon sampling theoremits various extensions and applications: A tutorial review", *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1565–1596, Nov. 1977, ISSN: 1558-2256. DOI: [10.1109/PROC.1977.10771](#).
 - [13] S. Wold, K. Esbensen, and P. Geladi, "Principal component analysis", *Chemometrics and Intelligent Laboratory Systems*, vol. 2, no. 1, pp. 37–52, Aug. 1, 1987, ISSN: 0169-7439. DOI: [10.1016/0169-7439\(87\)80084-9](#). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0169743987800849>.
 - [14] B. Iglewicz and D. C. Hoaglin, *How to detect and handle outliers*, ser. ASQC basic references in quality control v. 16. Milwaukee, Wis: ASQC Quality Press, 1993, 87 pp., ISBN: 978-0-87389-247-6.
 - [15] P. J. Rousseeuw and C. Croux, "Alternatives to the median absolute deviation", *Journal of the American Statistical Association*, vol. 88, no. 424, pp. 1273–1283, Dec. 1, 1993, ISSN: 0162-1459. DOI: [10.1080/01621459.1993.10476408](#). [Online]. Available: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1993.10476408>.
 - [16] D. Donoho, "De-noising by soft-thresholding", *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 613–627, May 1995, ISSN: 0018-9448, 1557-9654. DOI: [10.1109/18.382009](#).
 - [17] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage", *Journal of the American Statistical Association*, vol. 90, no. 432, pp. 1200–1224, Dec. 1, 1995, ISSN: 0162-1459. DOI: [10.1080/01621459.1995.10476626](#). [Online]. Available: <https://amstat.tandfonline.com/doi/abs/10.1080/01621459.1995.10476626>.
 - [18] A. Graps, "An introduction to wavelets", *IEEE Comput. Sci. Eng.*, vol. 2, no. 2, pp. 50–61, Jun. 1995, ISSN: 1070-9924. DOI: [10.1109/99.388960](#). [Online]. Available: <https://doi.org/10.1109/99.388960>.
 - [19] A. Schwarzenberg-Czerny, "On matrix factorization and efficient least squares solution.", *Astronomy and Astrophysics Supplement Series*, vol. 110, p. 405, Apr. 1, 1995, ISSN: 0365-0138. [Online]. Available: <http://adsabs.harvard.edu/abs/1995A%26AS..110..405S> (visited on 07/05/2020).
 - [20] A. van der Schaaf and J. H. van Hateren, "Modelling the power spectra of natural images: Statistics and information", *Vision Research*, vol. 36, no. 17, pp. 2759–2770, Sep. 1, 1996, ISSN: 0042-6989. DOI: [10.1016/0042-6989\(96\)00002-8](#). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0042698996000028>.
 - [21] D. B. Williams and C. B. Carter, "The transmission electron microscope", in *Transmission Electron Microscopy*, D. B. Williams and C. B. Carter, Eds., Boston, MA: Springer US, 1996, pp. 3–17, ISBN: 978-1-4757-2519-3. DOI: [10.1007/978-1-4757-2519-3_1](#). [Online]. Available: https://doi.org/10.1007/978-1-4757-2519-3_1.
 - [22] F. Guichard and F. Malgouyres, "Total variation based interpolation.", *Proc. European Signal Processing Conference*, vol. 3, Aug. 21, 1998.
 - [23] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration", *Medical Image Analysis*, vol. 2, no. 1, pp. 1–36, Mar. 1, 1998, ISSN: 1361-8415. DOI: [10.1016/S1361-8415\(01\)80026-8](#). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841501800268> (visited on 05/10/2019).
 - [24] R. A. Kruger, K. K. Kopecky, A. M. Aisen, D. R. Reinecke, G. A. Kruger, and W. L. Kiser, "Thermoacoustic CT with radio waves: A medical imaging paradigm", *Radiology*, vol. 211, no. 1, pp. 275–278, Apr. 1, 1999, ISSN: 0033-8419. DOI: [10.1148/](#)

- radiology.211.1.r99ap05275. [Online]. Available: <https://pubs.rsna.org/doi/abs/10.1148/radiology.211.1.r99ap05275>.
- [25] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization", *Nature*, vol. 401, no. 6755, p. 788, Oct. 1999, ISSN: 1476-4687. DOI: [10.1038/44565](https://doi.org/10.1038/44565). [Online]. Available: <https://www.nature.com/articles/44565>.
- [26] —, "Algorithms for non-negative matrix factorization", in *In NIPS*, MIT Press, 2000, pp. 556–562.
- [27] I. Pitas, *Digital Image Processing Algorithms and Applications*. John Wiley & Sons, Feb. 22, 2000, 436 pp., ISBN: 978-0-471-37739-9.
- [28] M. J. McAuliffe, F. M. Lalonde, D. McGarry, W. Gandler, K. Csaky, and B. L. Trus, "Medical image processing, analysis and visualization in clinical research", in *Proceedings 14th IEEE Symposium on Computer-Based Medical Systems. CBMS 2001*, Jul. 2001, pp. 381–386. DOI: [10.1109/CBMS.2001.941749](https://doi.org/10.1109/CBMS.2001.941749).
- [29] D. B. Murphy, *Fundamentals of light microscopy and electronic imaging*. New York: Wiley-Liss, 2001, 368 pp., ISBN: 978-0-471-25391-4.
- [30] I. N. Aizenberg and C. Butakoff, "Frequency domain medianlike filter for periodic and quasi-periodic noise removal", in *Image Processing: Algorithms and Systems*, vol. 4667, International Society for Optics and Photonics, May 22, 2002, pp. 181–192. DOI: [10.1117/12.467980](https://doi.org/10.1117/12.467980). [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/4667/0000/Frequency-domain-medianlike-filter-for-periodic-and-quasi-periodic-noise/10.1117/12.467980.short>.
- [31] I. T. Jolliffe, *Principal Component Analysis*, 2nd edition. New York: Springer, Oct. 1, 2002, 488 pp., ISBN: 978-0-387-95442-4.
- [32] D. N. Sidorov and A. C. Kokaram, "Suppression of moire patterns via spectral analysis", presented at the Electronic Imaging 2002, C.-C. J. Kuo, Ed., San Jose, CA: SPIE Press, Jan. 7, 2002, p. 895. DOI: [10.1117/12.453134](https://doi.org/10.1117/12.453134). [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.453134>.
- [33] A. Pizurica, W. Philips, I. Lemahieu, and M. Achero, "A versatile wavelet domain noise filtration technique for medical imaging", *IEEE Transactions on Medical Imaging*, vol. 22, no. 3, pp. 323–331, Mar. 2003, ISSN: 0278-0062. DOI: [10.1109/TMI.2003.809588](https://doi.org/10.1109/TMI.2003.809588). [Online]. Available: <http://ieeexplore.ieee.org/document/1199634/> (visited on 01/18/2019).
- [34] K. Downard, *Mass Spectrometry*. Royal Society of Chemistry, Aug. 5, 2004, ISBN: 978-0-85404-609-6. DOI: [10.1039/9781847551306](https://doi.org/10.1039/9781847551306). [Online]. Available: <https://pubs.rsc.org/en/content/ebook/978-0-85404-609-6>.
- [35] L. J. Erasmus, D. Hurter, M. Naude, H. G. Kritzing, and S. Acho, "A short overview of MRI artefacts", *South African Journal of Radiology*, vol. 8, no. 2, p. 13, Jun. 9, 2004, ISSN: 2078-6778. DOI: [10.4102/sajr.v8i2.127](https://doi.org/10.4102/sajr.v8i2.127). [Online]. Available: <https://sajr.org.za/index.php/sajr/article/view/127>.
- [36] P. Gravel, G. Beaudoin, and J. DeGuisé, "A method for modeling noise in medical images", *IEEE Transactions on Medical Imaging*, vol. 23, no. 10, pp. 1221–1232, Oct. 2004, ISSN: 0278-0062. DOI: [10.1109/TMI.2004.832656](https://doi.org/10.1109/TMI.2004.832656). [Online]. Available: <http://ieeexplore.ieee.org/document/1339429/> (visited on 01/18/2019).
- [37] P. O. Hoyer, "Non-negative matrix factorization with sparseness constraints", *arXiv:cs/0408058*, Aug. 25, 2004. arXiv: [cs/0408058](https://arxiv.org/abs/cs/0408058). [Online]. Available: <http://arxiv.org/abs/cs/0408058>.
- [38] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: From error visibility to structural similarity", *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004, ISSN: 1057-7149, 1941-0042. DOI: [10.1109/TIP.2003.819861](https://doi.org/10.1109/TIP.2003.819861).

- [39] A. Buades, B. Coll, and J. M. Morel, “A review of image denoising algorithms, with a new one”, *Multiscale Modeling & Simulation*, vol. 4, no. 2, pp. 490–530, Jan. 2005, ISSN: 1540-3459, 1540-3467. DOI: [10.1137/040616024](https://doi.org/10.1137/040616024). [Online]. Available: <http://epubs.siam.org/doi/10.1137/040616024>.
- [40] R. L. Lagendijk, P. M. B. van Roosmalen, J. Biemond, A. Rare, and M. J. T. Rein-
ders, “3.11 - video enhancement and restoration”, in *Handbook of Image and Video Processing (Second
ser. Communications, Networking and Multimedia*, A. Bovik, Ed., Burlington: Aca-
demic Press, Jan. 1, 2005, pp. 275–VI, ISBN: 978-0-12-119792-6. DOI: [10.1016/B978-
012119792-6/50080-2](https://doi.org/10.1016/B978-012119792-6/50080-2). [Online]. Available: [http://www.sciencedirect.com/
science/article/pii/B9780121197926500802](http://www.sciencedirect.com/science/article/pii/B9780121197926500802).
- [41] E. Grafarend, “Linear and nonlinear models: Fixed effects, random effects, and mixed
models”, Jan. 1, 2006.
- [42] R. Maronna, D. Martin, and V. Yohai, “Robust statistics: Theory and methods”,
Jun. 1, 2006. DOI: [10.1002/0470010940](https://doi.org/10.1002/0470010940).
- [43] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince, *MRI from Picture to Proton*,
2nd ed. Cambridge: Cambridge University Press, 2006, ISBN: 978-0-511-54540-5. DOI:
[10.1017/CB09780511545405](https://doi.org/10.1017/CB09780511545405). [Online]. Available: [https://www.cambridge.org/
core/books/mri-from-picture-to-proton/3ADC814FF8FC6A78A54D37746F806D5A](https://www.cambridge.org/core/books/mri-from-picture-to-proton/3ADC814FF8FC6A78A54D37746F806D5A)
(visited on 07/04/2020).
- [44] F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David, “Phase retrieval and differential
phase-contrast imaging with low-brilliance x-ray sources”, *Nature Physics*, vol. 2,
no. 4, p. 258, Apr. 2006, ISSN: 1745-2481. DOI: [10.1038/nphys265](https://doi.org/10.1038/nphys265). [Online]. Avail-
able: <https://www.nature.com/articles/nphys265> (visited on 06/14/2019).
- [45] A. Pizurica, A. M. Wink, E. Vansteenkiste, and W. P. B. J. T. M. Roerdink. (May 1,
2006). A review of wavelet denoising in MRI and ultrasound brain imaging, *Current
Medical Imaging*, [Online]. Available: [http://www.eurekaselect.com/55848/
article](http://www.eurekaselect.com/55848/article) (visited on 08/01/2019).
- [46] Y. Wang and H. Zhou, “Total variation wavelet-based medical image denoising”,
International Journal of Biomedical Imaging, vol. 2006, pp. 1–6, 2006, ISSN: 1687-
4188, 1687-4196. DOI: [10.1155/IJBI/2006/89095](https://doi.org/10.1155/IJBI/2006/89095). [Online]. Available: [http://
www.hindawi.com/journals/ijbi/2006/089095/abs/](http://www.hindawi.com/journals/ijbi/2006/089095/abs/) (visited on 01/18/2019).
- [47] L. A. McDonnell and R. M. A. Heeren, “Imaging mass spectrometry”, *Mass Spectrometry Reviews*,
vol. 26, no. 4, pp. 606–643, 2007, ISSN: 1098-2787. DOI: [10.1002/mas.20124](https://doi.org/10.1002/mas.20124). [Online].
Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mas.20124>.
- [48] S. Singh and A. Goyal, “The origin of echocardiography”, *Texas Heart Institute Journal*,
vol. 34, no. 4, pp. 431–438, 2007, ISSN: 0730-2347. [Online]. Available: [https://www.
ncbi.nlm.nih.gov/pmc/articles/PMC2170493/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2170493/).
- [49] I. Aizenberg and C. Butakoff, “A windowed gaussian notch filter for quasi-periodic
noise removal”, *Image and Vision Computing*, vol. 26, no. 10, pp. 1347–1353, Oct. 1,
2008, ISSN: 0262-8856. DOI: [10.1016/j.imavis.2007.08.011](https://doi.org/10.1016/j.imavis.2007.08.011). [Online]. Available:
<http://www.sciencedirect.com/science/article/pii/S0262885607001382>.
- [50] P. W. Siy, R. A. Moffitt, R. M. Parry, Y. Chen, Y. Liu, M. C. Sullards, A. H. Merrill,
and M. D. Wang, “Matrix factorization techniques for analysis of imaging mass spec-
trometry data”, in *2008 8th IEEE International Conference on BioInformatics and BioEngineering*,
Athens, Greece: IEEE, Oct. 2008, pp. 1–6, ISBN: 978-1-4244-2844-1. DOI: [10.1109/
BIBE.2008.4696797](https://doi.org/10.1109/BIBE.2008.4696797). [Online]. Available: [http://ieeexplore.ieee.org/document/
4696797/](http://ieeexplore.ieee.org/document/4696797/) (visited on 01/18/2019).
- [51] R. P. Erni, “Atomic resolution imaging with a sub-50 pm electron probe”, *Lawrence Berkeley National L*
Jul. 22, 2009. [Online]. Available: <https://escholarship.org/uc/item/3cs0m4vr>
(visited on 07/20/2019).

- [52] N. P. Hurley and S. T. Rickard, “Comparing measures of sparsity”, arXiv:0811.4706 [cs, math], Apr. 27, 2009. arXiv: [0811.4706](https://arxiv.org/abs/0811.4706). [Online]. Available: <http://arxiv.org/abs/0811.4706>.
- [53] R. B. Paranjape, “Chapter 1 - fundamental enhancement techniques”, in *Handbook of Medical Image Processing*, I. N. Bankman, Ed., Burlington: Academic Press, Jan. 1, 2009, pp. 3–18, ISBN: 978-0-12-373904-9. DOI: [10.1016/B978-012373904-9.50008-8](https://doi.org/10.1016/B978-012373904-9.50008-8). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780123739049500088>.
- [54] J. Ayache, L. Beaunier, J. Boumendil, G. Ehret, and D. Laub, “Artifacts in transmission electron microscopy”, in *Sample Preparation Handbook for Transmission Electron Microscopy: Methods and Techniques*, J. Ayache, L. Beaunier, J. Boumendil, G. Ehret, and D. Laub, Eds., New York, NY: Springer New York, 2010, pp. 125–170, ISBN: 978-0-387-98182-6. DOI: [10.1007/978-0-387-98182-6_6](https://doi.org/10.1007/978-0-387-98182-6_6). [Online]. Available: https://doi.org/10.1007/978-0-387-98182-6_6.
- [55] C. Chaux, A. Benazza-Benyahia, J.-C. Pesquet, and L. Duval, “Chapitre 7 - wavelet transform for the denoising of multivariate images”, in *Wavelet Transform and its Applications*, Jan. 1, 2010, ISBN: 978-1-84821-139-1.
- [56] J. M. Bini, J. Spain, K. S. Nehal, M. Rajadhyaksha, V. Hazelwood, and C. A. Di-Marzio, “Confocal mosaicing microscopy of human skin ex vivo spectral analysis for digital staining to simulate histology-like appearance”, *Journal of Biomedical Optics*, vol. 16, no. 7, p. 076008, Jul. 2011, ISSN: 1083-3668, 1560-2281. DOI: [10.1117/1.3596742](https://doi.org/10.1117/1.3596742). [Online]. Available: <https://www.spiedigitallibrary.org/journals/Journal-of-Biomedical-Optics/volume-16/issue-7/076008/----Custom-HTML---Confocal/10.1117/1.3596742.short>.
- [57] J. T. Bushberg and J. M. Boone, *The Essential Physics of Medical Imaging*. Lippincott Williams & Wilkins, Dec. 20, 2011, 1049 pp., Google-Books-ID: tqM8IG3f8bsC, ISBN: 978-0-7817-8057-5.
- [58] E. J. Candes, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?”, *Journal of the ACM*, vol. 58, no. 3, p. 39, May 2011, ISSN: 0004-5411. DOI: [10.1145/1970392.1970395](https://doi.org/10.1145/1970392.1970395).
- [59] R. Javan, “Fundamentals behind the 10 most common magnetic resonance imaging artifacts with correction strategies and 10 high-yield points”, *European Congress of Radiology*, 2011. DOI: [10.1594/ecr2011/c-1248](https://doi.org/10.1594/ecr2011/c-1248). [Online]. Available: <http://epos.myesr.org/poster/ecr2011/C-1248>.
- [60] B. Ergen, “Signal and image denoising using wavelet transform”, in *Advances in Wavelet Theory and T*, InTech, Apr. 4, 2012, ISBN: 978-953-51-0494-0. DOI: [10.5772/36434](https://doi.org/10.5772/36434). [Online]. Available: <http://www.intechopen.com/books/advances-in-wavelet-theory-and-their-applications-in-engineering-physics-and-technology/wavelet-signal-and-image-denoising>.
- [61] A. Mittal, A. K. Moorthy, and A. C. Bovik, “No-reference image quality assessment in the spatial domain”, *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4695–4708, Dec. 2012, Conference Name: IEEE Transactions on Image Processing, ISSN: 1941-0042. DOI: [10.1109/TIP.2012.2214050](https://doi.org/10.1109/TIP.2012.2214050).
- [62] Z. Wei, J. Wang, H. Nichol, S. Wiebe, and D. Chapman, “A median-gaussian filtering framework for moiré pattern noise removal from x-ray microscopy image”, *Micron (Oxford, England : 1993)*, vol. 43, no. 2, pp. 170–176, Feb. 2012, ISSN: 0968-4328. DOI: [10.1016/j.micron.2011.07.009](https://doi.org/10.1016/j.micron.2011.07.009). [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3858302/>.
- [63] S. Abeytunge, Y. Li, B. A. Larson, G. Peterson, E. Seltzer, R. Toledo-Crow, and M. Rajadhyaksha, “Confocal microscopy with strip mosaicing for rapid imaging over large areas of excised tissue”, *Journal of Biomedical Optics*, vol. 18, no. 6, p. 061227, Feb. 2013, ISSN: 1083-3668, 1560-2281. DOI: [10.1117/1.JBO.18.6.061227](https://doi.org/10.1117/1.JBO.18.6.061227). [On-

- line]. Available: <https://www.spiedigitallibrary.org/journals/Journal-of-Biomedical-Optics/volume-18/issue-6/061227/Confocal-microscopy-with-strip-mosaicing-for-rapid-imaging-over-large/10.1117/1.JBO.18.6.061227.short>.
- [64] I. Amidror, Mastering the Discrete Fourier Transform in One, Two or Several Dimensions, ser. Computational Imaging and Vision. London: Springer London, 2013, vol. 43, ISBN: 978-1-4471-5166-1 978-1-4471-5167-8. DOI: [10.1007/978-1-4471-5167-8](https://doi.org/10.1007/978-1-4471-5167-8). [Online]. Available: <http://link.springer.com/10.1007/978-1-4471-5167-8>.
- [65] L. Tan and J. Jiang, “Chapter 14 - image processing basics”, in Digital Signal Processing (Second Edition) L. Tan and J. Jiang, Eds., Boston: Academic Press, Jan. 1, 2013, pp. 683–765, ISBN: 978-0-12-415893-1. DOI: [10.1016/B978-0-12-415893-1.00014-7](https://doi.org/10.1016/B978-0-12-415893-1.00014-7). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780124158931000147>.
- [66] J. K. C. Chan, “The wonderful colors of the hematoxylineosin stain in diagnostic surgical pathology”, International Journal of Surgical Pathology, vol. 22, no. 1, pp. 12–32, Feb. 1, 2014, ISSN: 1066-8969. DOI: [10.1177/1066896913517939](https://doi.org/10.1177/1066896913517939). [Online]. Available: <https://doi.org/10.1177/1066896913517939>.
- [67] S. Chatterjee, “Artefacts in histopathology”, Journal of Oral and Maxillofacial Pathology, vol. 18, no. 4, p. 111, Sep. 1, 2014, ISSN: 0973-029X. DOI: [10.4103/0973-029X.141346](https://doi.org/10.4103/0973-029X.141346). [Online]. Available: <http://www.jomfp.in/article.asp?issn=0973-029X;year=2014;volume=18;issue=4;spage=111;epage=116;aulast=Chatterjee;type=0>.
- [68] Z. Liu, L. Tian, S. Liu, and L. Waller, “Real-time brightfield, darkfield, and phase contrast imaging in a light-emitting diode array microscope”, Journal of Biomedical Optics, vol. 19, no. 10, Oct. 2014, ISSN: 1083-3668, 1560-2281. DOI: [10.1117/1.JBO.19.10.106002](https://doi.org/10.1117/1.JBO.19.10.106002). [Online]. Available: <https://www.spiedigitallibrary.org/journals/Journal-of-Biomedical-Optics/volume-19/issue-10/106002/Real-time-brightfield-darkfield-and-phase-contrast-imaging-in-a/10.1117/1.JBO.19.10.106002.short>.
- [69] A. K. Boyat and B. K. Joshi, “A review paper: Noise models in digital image processing”, Signal & Image Processing : An International Journal (SIPIJ), vol. 6, no. 2, May 13, 2015. arXiv: [1505.03489](https://arxiv.org/abs/1505.03489). [Online]. Available: <http://arxiv.org/abs/1505.03489>.
- [70] G.-M. Ionita, D. Coltuc, S. G. Stanciu, and D. E. Tranca, “Automatic moiré pattern removal in microscopic images”, in 2015 19th International Conference on System Theory, Control and Cheile Gradistei, Romania: IEEE, Oct. 2015, pp. 776–779, ISBN: 978-1-4799-8481-7. DOI: [10.1109/ICSTCC.2015.7321388](https://doi.org/10.1109/ICSTCC.2015.7321388). [Online]. Available: <http://ieeexplore.ieee.org/document/7321388/>.
- [71] I. Irum, M. A. Shahid, M. Sharif, and M. Raza, “A review of image denoising methods”, Journal of Engineering Science and Technology Review, p. 9, 2015.
- [72] F. Sur, “An a-contrario approach to quasi-periodic noise removal”, in 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Quebec City, QC, Canada: IEEE, Sep. 2015, pp. 3841–3845, ISBN: 978-1-4799-8339-1. DOI: [10.1109/ICIP.2015.7351524](https://doi.org/10.1109/ICIP.2015.7351524). [Online]. Available: <http://ieeexplore.ieee.org/document/7351524/>.
- [73] —, “Total variation minimization for spectrum interpolation in quasiperiodic noise removal”, Université de Lorraine, p. 6, Jun. 2015. [Online]. Available: https://members.loria.fr/FSur/software/ACARPENOS/note_TVmin.pdf.
- [74] F. Sur and M. Grediac, “Automated removal of quasiperiodic noise using frequency domain statistics”, Journal of Electronic Imaging, vol. 24, no. 1, p. 013003, Feb. 2015, ISSN: 1017-9909, 1560-229X. DOI: [10.1117/1.JEI.24.1.013003](https://doi.org/10.1117/1.JEI.24.1.013003). [Online]. Available: <https://www.spiedigitallibrary.org/journals/Journal-of-Electronic-Imaging/volume-24/issue-1/013003/Automated-removal-of->

- [quasiperiodic-noise-using-frequency-domain-statistics/10.1117/1.JEI.24.1.013003.short](#).
- [75] F. Sur and M. Grédiac, “Automated removal of quasiperiodic noise using frequency domain statistics”, *Journal of Electronic Imaging*, vol. 24, no. 1, p. 013 003, Feb. 11, 2015, ISSN: 1017-9909. DOI: [10.1117/1.JEI.24.1.013003](#). [Online]. Available: <http://electronicimaging.spiedigitallibrary.org/article.aspx?doi=10.1117/1.JEI.24.1.013003> (visited on 01/18/2019).
 - [76] L. Gondara, “Medical image denoising using convolutional denoising autoencoders”, in *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, Barcelona, Spain: IEEE, Dec. 2016, pp. 241–246, ISBN: 978-1-5090-5910-2. DOI: [10.1109/ICDMW.2016.0041](#). [Online]. Available: <http://ieeexplore.ieee.org/document/7836672/> (visited on 01/18/2019).
 - [77] M. Grediac, F. Sur, and B. Blaysat, “Removing quasi-periodic noise in strain maps by filtering in the fourier domain”, *Experimental Techniques*, vol. 40, no. 3, pp. 959–971, 2016. DOI: [10.1007/s40799-016-0100-2](#). [Online]. Available: <https://hal.inria.fr/hal-01163838/document> (visited on 02/10/2019).
 - [78] I. T. Jolliffe and J. Cadima, “Principal component analysis: A review and recent developments”, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 374, no. 2065, p. 20 150 202, Apr. 13, 2016, Publisher: Royal Society. DOI: [10.1098/rsta.2015.0202](#). [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsta.2015.0202> (visited on 07/05/2020).
 - [79] P. Källback, A. Nilsson, M. Shariatgorji, and P. E. Andrén, “msIQuant quantitation software for mass spectrometry imaging enabling fast access, visualization, and analysis of large data sets”, *Analytical Chemistry*, vol. 88, no. 8, pp. 4346–4353, Apr. 19, 2016, ISSN: 0003-2700. DOI: [10.1021/acs.analchem.5b04603](#). [Online]. Available: <https://doi.org/10.1021/acs.analchem.5b04603>.
 - [80] J. C. Russ, *The Image Processing Handbook*. CRC Press, Apr. 19, 2016, ISBN: 978-0-429-10601-9. DOI: [10.1201/b10720](#). [Online]. Available: <https://www.taylorfrancis.com/books/9780429106019>.
 - [81] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 4 edition. New York, NY: Pearson, Mar. 30, 2017, 1192 pp., ISBN: 978-0-13-335672-4.
 - [82] M. G. Ionita and H. G. Coanda, “Automatic periodic noise removal in microscopy images”, in *2017 International Symposium on Signals, Circuits and Systems (ISSCS)*, Iasi, Romania: IEEE, Jul. 2017, pp. 1–4, ISBN: 978-1-5386-0674-2. DOI: [10.1109/ISSCS.2017.8034894](#). [Online]. Available: <http://ieeexplore.ieee.org/document/8034894/>.
 - [83] (Aug. 2, 2017). Sample preparation for fluorescence microscopy, News-Medical.net, [Online]. Available: <https://www.news-medical.net/life-sciences/Sample-Preparation-for-Fluorescence-Microscopy.aspx> (visited on 08/09/2019).
 - [84] M. Ionita and H. Coanda, “Wavelet and fourier decomposition based periodic noise removal in microscopy images”, *The Scientific Bulletin of Electrical Engineering Faculty*, vol. 18, no. 1, pp. 68–71, Apr. 1, 2018, ISSN: 2286-2455. DOI: [10.1515/sbeef-2017-0025](#). [Online]. Available: <http://content.sciendo.com/view/journals/sbeef/18/1/article-p68.xml>.
 - [85] D. NASA Inductiveload, *English: Electromagnetic spectrum properties include wavelength, frequency and amplitude*, Jan. 3, 2019. [Online]. Available: https://commons.wikimedia.org/wiki/File:Electromagnetic_spectrum_with_sources.svg (visited on 11/22/2019).
 - [86] I. Amidror, *The theory of the Moiré phenomenon*, 2nd ed, ser. Computational imaging and vision. Dordrecht: Kluwer Academic, vol. 38, 529 pp., ISBN: 978-1-84882-180-4.

- [87] (). Bedbug | free stock photo | microscopic bedbug scanning electron micrograph. | #16714, [Online]. Available: <http://www.freestockphotos.biz/stockphoto/16714> (visited on 11/22/2019).
- [88] C. Croux, P. Filzmoser, and M. R. Oliveira, “Algorithms for projection-pursuit robust principal component analysis”, p. 19,
- [89] P. Paul, D. Kalamatianos, H. Duessmann, and H. Huber, Automatic quality assessment for fluorescence ISBN: 978-1-4244-2844-1 978-1-4244-2845-8. DOI: [10.1109/BIBE.2008.4696665](https://doi.org/10.1109/BIBE.2008.4696665). [Online]. Available: <https://www.infona.pl/resource/bwmeta1.element.ieee-art-000004696665>.
- [90] (). Principal component analysis explained visually, Explained Visually, [Online]. Available: <http://setosa.io/ev/principal-component-analysis/> (visited on 09/19/2019).
- [91] (). Quality analysis archives, Eric Hanson, [Online]. Available: <http://ericjohnhanson.com/tag/quality-analysis/> (visited on 08/09/2019).
- [92] E. M. S. Springer Verlag GmbH. (). Encyclopedia of mathematics, [Online]. Available: https://www.encyclopediaofmath.org/index.php/Main_Page (visited on 07/25/2019).
- [93] U. R. V. (). Noise | radiology reference article | radiopaedia.org, Radiopaedia, [Online]. Available: <https://radiopaedia.org/articles/noise> (visited on 08/02/2019).
- [94] L. J. van Vliet, F. R. Boddeke, D. Sudar, and I. T. Young, “Image detectors for digital image microscopy”, pp. 37–64,