UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM COMPUTAÇÃO

RAFAEL LUCENA GERMANO

**Real-Time Frequency Adjustment of
Images and Videos**

Thesis presented in partial fulfillment of the
requirements for the degree of Master of Computer
Science

Advisor: Prof. Dr. Eduardo Simões Lopes Gastal

Porto Alegre
July 2021

CIP — CATALOGING-IN-PUBLICATION

Germano, Rafael Lucena

Real-Time Frequency Adjustment of Images and Videos /
Rafael Lucena Germano. – Porto Alegre: PPGC da UFRGS, 2021.

62 f.: il.

Thesis (Master) – Universidade Federal do Rio Grande do
Sul. Programa de Pós-Graduação em Computação, Porto Alegre,
BR–RS, 2021. Advisor: Eduardo Simões Lopes Gastal.

1. Image Processing. I. Gastal, Eduardo Simões Lopes. II. Tí-
tulo.

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

Reitor: Prof. Carlos André Bulhões

Vice-Reitora: Prof^a. Patricia Pranke

Pró-Reitor de Pós-Graduação: Prof. Celso Giannetti Loureiro Chaves

Diretora do Instituto de Informática: Prof^a. Carla Maria Dal Sasso Freitas

Coordenador do PPGC: Prof. Claudio Rosito Jung

Bibliotecária-chefe do Instituto de Informática: Beatriz Regina Bastos Haro

ABSTRACT

We present a technique for real-time adjustment of spatial frequencies in images and videos. Our method allows for both decreasing and increasing of frequencies, and is orthogonal to image resizing. Thus, it can be used to automatically adjust spatial frequencies to preserve the appearance of structured patterns during image downscaling and upscaling. By pre-computing the image's space-frequency decomposition and its unwrapped phases, these operations can be performed in real time, thanks to our novel mathematical perspective on frequency manipulation of digital images: interpreting the problem through the theory of instantaneous frequencies and phase unwrapping. To make this possible, we introduce an algorithm for the simultaneous phase unwrapping of several unordered frequency components, which also deals with the frequency-sign ambiguity of real signals. As such, our method provides theoretical and practical improvements to the concept of spectral remapping, enabling real-time performance and improved color handling. We demonstrate its effectiveness on a large number of images subject to frequency adjustment. By providing real-time control over the spatial frequencies associated with structured patterns, our technique expands the range of creative and technical possibilities for image and video processing.

Keywords: Image and Video Processing, Frequency Analysis, Frequency Adjustment.

Ajuste de Frequência em Tempo Real de Imagens e Vídeos

RESUMO EXPANDIDO EM PORTUGUÊS

Apresentamos uma técnica para ajuste em tempo real de frequências espaciais em imagens e vídeos. Nosso método permite aumentar e diminuir as frequências, e é ortogonal ao redimensionamento de imagens. Portanto, pode ser utilizado para automaticamente ajustar frequências espaciais preservando a aparência de padrões estruturados durante a redução e ampliação de imagens. Pré-computando a decomposição espaço-frequência da imagem e sua *unwrapped phase*, essas operações podem ser feitas em tempo real, devido à nossa nova perspectiva matemática sobre manipulação de frequências em imagens digitais: interpretando o problema por meio da teoria de frequências instantâneas e phase unwrapping. Para esse fim, introduzimos um algoritmo de phase unwrapping simultâneo de vários componentes de frequência desordenados, que também lida com a ambiguidade de sinal em sinais reais. Desse modo, nosso método proporciona melhorias teóricas e práticas para o conceito de remapeamento espectral, permitindo desempenho em tempo real e melhor tratamento de cores. Demonstramos sua eficácia em um grande número de imagens sujeitas a ajustes de frequência. Ao proporcionar controle em tempo real de frequências espaciais associadas a padrões estruturados, nosso método expande uma gama de possibilidades criativas e técnicas para processamento de imagens e vídeos.

Nosso método inicialmente detecta as *frequências e amplitudes* de várias ondas senoidais que localmente reconstroem uma imagem por meio de uma transformada de Gabor e remove as ondas da imagem para obter o *resíduo*. Em seguida, para obter um *campo de fases suave*, as frequências são integradas usando o nosso método de integração por mínimos quadrados ponderado por uma medida de similaridade. Finalmente, a imagem é reconstruída somando o resíduo da imagem original e o produto das amplitudes e os cossenos das fases escaladas por um fator de ajuste. Além disso, podemos alterar as frequências em tempo real repetindo o processo de reconstrução com outro fator de ajuste. Demonstramos a performance em tempo real com uma implementação em GPU, nossa implementação usa a API OpenGL e o fato de que a reconstrução da imagem usa operações por pixel para reconstruir a imagem usando o fragment shader do pipeline gráfico.

Palavras-chave: Processamento de Imagens e Vídeos, Análise de Frequência, Ajuste de Frequência.

LIST OF ABBREVIATIONS AND ACRONYMS

DFT	Discrete Fourier Transform
HHT	Hilbert-Huang Transform
IDFT	Inverse Discrete Fourier Transform
IMF	Intrinsic Mode Function
O-MOMS	Optimal Maximal-Order-Minimal-Support
PCA	Principal Component Analysis
SR	Spectral Remapping

LIST OF SYMBOLS

$\alpha \in \mathbb{R}$	Adjustment factor
$\theta^n(p) \in \mathbb{R}$	Analytic phase of the n -th wave at pixel p
$\nabla\theta^n(p) \in [-\pi, \pi]^d$	Analytic Instantaneous Frequency of the n -th wave at pixel p in an d -dimensional signal ($d = 2$ for images and $d = 3$ for videos)
$\kappa_p(n) \in \mathbb{N}$	Permutation of the n -th wave at pixel p into the $\kappa_p(n)$ -th <i>detected</i> wave at pixel p
$\hat{\varphi}^n(p) \in \mathbb{R}$	Wrapped phase of the n -th <i>detected</i> wave at pixel p
$\varphi^n(p) \in \mathbb{R}$	Unwrapped phase of the n -th <i>detected</i> wave at pixel p
$\nabla\varphi^n(p) \in [-\pi, \pi]^d$	Instantaneous Frequency of the n -th <i>detected</i> wave at pixel p in an d -dimensional signal
$\sigma^n(p) \in \{-1, +1\}$	Sign function associated to the phase of the n -th <i>detected</i> wave at pixel p
$\mu_{pq}^{nm} \in [0, 1]$	Similarity measure between the n -th <i>detected</i> wave of pixel p and the m -th <i>detected</i> wave of pixel q
$A^n(p) \in \mathbb{R}^+$	Analytic amplitude of the n -th wave at pixel p
$B^n(p) \in \mathbb{R}^+$	Amplitude of the n -th <i>detected</i> wave at pixel p
$T_x/T_y/T_\circ$	Unit-shift operators
$\mathcal{E} : \mathbb{R}^N \rightarrow \mathbb{R}$	Energy functional for an image containing N waves in total
$\chi_{c_1 c_2} : \mathbb{R}^N \rightarrow \mathbb{R}$	Cross-channel phase constraint functional between channels c_1 and c_2 , for an image containing N waves in total
$\eta_{c_1 c_2}^{nm}(p) \in \mathbb{R}$	Similarity measure between the frequency vectors of the n -th and m -th <i>detected</i> waves in channels c_1 and c_2 of pixel p

LIST OF FIGURES

Figure 1.1	Example of frequency adjustments applied to “Brothers”	10
Figure 1.2	Overview of the frequency-adjustment technique	12
Figure 2.1	Examples of downscaling	16
Figure 2.2	Example of windowing used for frequency detection	19
Figure 3.1	Frequency adjustment of a Perlin noise scalar field	25
Figure 3.2	Unsuccessful frequency adjustment of a Perlin noise scalar field	26
Figure 4.1	Example of multi-channel phase unwrapping	34
Figure 5.1	Decomposition used for frequency adjustment	42
Figure 5.2	Perceptual frequency adjustment	43
Figure 5.3	Anisotropic frequency adjustment of a radial chirp	44
Figure 5.4	Anisotropic frequency adjustment of barbara.....	45
Figure 6.1	Frequency adjustment of complex natural patterns	47
Figure 6.2	Example of localized frequency adjustment.....	48
Figure 6.3	Examples of downscaling and upscaling with frequency adjustments.....	50
Figure 6.4	Picture of a lizard downscaled with our frequency adjustment technique	52
Figure 6.5	Example of frequency adjustment applied to a shirt	53
Figure 6.6	Example of video frequency adjustment	54
Figure 7.1	Chromatic distortions from multi-channel alignment	56

CONTENTS

1 INTRODUCTION.....	9
2 RELATED WORKS AND TECHNICAL BACKGROUND	13
2.1 Image resizing.....	14
2.1.1 Overview on Spectral Remapping	15
2.2 Technical Background	15
2.2.1 The discrete Fourier transform.....	17
2.2.2 Space-frequency decomposition: Wave detection	18
2.2.3 Phase Unwrapping	19
2.2.3.1 Unwrapping the phase of one complex wave	20
2.2.3.2 Common approach for estimating the wrapped gradient field for one com- plex wave	21
2.2.4 Estimating the phase gradient field for one real wave	22
3 FREQUENCY ADJUSTMENT.....	24
3.1 Frequency Adjustment of Natural Images	24
3.2 Finding a Decomposition.....	25
4 PHASE UNWRAPPING IN TWO DIMENSIONS	27
4.1 Unwrapping the phase of one real wave	27
4.1.1 Integrating the phase gradient field for one real wave.	28
4.1.2 Determining sign changes and computing $\Delta\sigma$	30
4.2 Unwrapping the phases of several real waves	31
4.3 Processing Color Images	32
5 IMPLEMENTATION DETAILS.....	36
5.1 Sparse Sampling.....	36
5.2 Use with Image Resizing.....	37
5.2.1 Choosing α	38
5.3 Processing Videos and 3-D Volumes.....	39
5.4 Splitted phase unwrapping for anisotropic frequency adjustment	40
5.5 Complexity Analysis	41
6 RESULTS.....	46
7 THEORETICAL DISCUSSION	55
7.1 Limitations.....	55
7.2 Comparison to Spectral Remapping	56
7.3 Future Works.....	57
8 CONCLUSION	59
REFERENCES.....	60

1 INTRODUCTION

The frequency domain is a common way of analyzing digital signals. It is used in digital signal processing as a tool for filtering, compression, and several other applications (OPPENHEIM; SCHAFER, 2010). For some applications, however, *time*-frequency analysis is an alternative and preferred tool, since it allows for the study of transient phenomena. In image and video processing, the analogous *space*-frequency domain has also been extensively used for example with wavelets for image compression, data analysis and denoising (PRASAD; IYENGAR, 1997; BARBHUIYA; HEMACHANDRAN, 2013). Space-frequency analysis may be preferred to traditional frequency analysis because the signal becomes represented by waves concentrated both in space and frequency (MALLAT, 2009). This is useful because the frequency content becomes more meaningful in each associated region.

We present a new mathematical perspective on frequency manipulation of digital images and videos. Using the theory of instantaneous frequencies and phase unwrapping operators (MALLAT, 2009), our **real-time spatial frequency adjustment technique** is the first method capable of performing image edits such as the ones shown in Figure 1.1 in real time. The performance improvement resulting from our formulation is also critical for making frequency adjustments in videos practical. The idea of our method is to decompose an image into local phase and amplitude components such that the phase encodes the high-frequency structured content. This decomposition provides simple and interesting ways of manipulating frequency through phase manipulation.

Our technique starts by obtaining a space-frequency decomposition of the image (**orange block** of Figure 1.2). The decomposition represents the original image as a sum of several cosine functions and a residual. Each cosine has an associated frequency and amplitude. The frequencies are jointly integrated by our weighted least-squares system (**blue block** of Figure 1.2) to obtain smooth phase functions that may be scaled by a factor α to reconstruct the image with the original frequencies adjusted by α (**green block** of Figure 1.2). The reconstruction algorithm may be performed per-pixel, allowing for a fast parallel real-time implementation. Despite working with cosine functions for the decomposition, our technique is able to handle complex structured patterns as demonstrated throughout this thesis (e.g., see Figure 6.1).

Our novel mathematical framework is fundamental in allowing our approach to work in real time. To make this possible, we introduce technical contributions to ridge

Figure 1.1 – (a) Reference image with $1,618 \times 2,048$ pixels. (b–d) Examples of frequency adjustments applied to (a) using our technique, with the images kept at the original resolution. (e) Downscaled version of (a) generated with our technique in real time (0.01 seconds for frequency adjustment and image reconstruction). (f) Downscaled version of (a) obtained with spectral remapping (SR) (GASTAL; OLIVEIRA, 2017) (1.5 seconds). (g) Downscaled version of (a) without changing the image frequencies. In images (e–g), resampling to 135×171 pixels was done using the O-MOMS method described by Nehab and Hoppe (NEHAB; HOPPE, 2014).

Source: (a) Hand drawing “Brothers” courtesy of Tyler Hobbs. (b–g) The authors.

analysis and 2-D/3-D phase unwrapping of digital images and videos (Section 4.2). These contributions are essential for handling complex patterns such as the sunflower’s disk in Figure 6.1, or the lizard’s leg in Figure 6.4. We also show how to handle the problematic sign-ambiguity that arises from the conjugate-symmetry of real waves (Section 4.1.1). Our approach also improves color handling when compared to recent frequency remapping methods (GASTAL; OLIVEIRA, 2017), by correlating the color channels during frequency integration (Section 4.3). This improvement is visible in Figure 4.1.

By defining the frequency scaling factor α as a function of an image resizing factor r , spatial frequencies are automatically adjusted to preserve the appearance of structured patterns during image downscaling (Section 5.2.1, Figure 1.1(e)). Combined with our method’s real-time performance, this opens up a variety of interesting applications, such as dynamically computing detail-preserving thumbnails for image galleries, or adapting an image’s content in real time based on viewing distance. We also demonstrate the use of our technique for frequency upscaling (Figure 6.5), and how it can be combined with existing super-resolution techniques to obtain novel super-resolution results (Figure 6.3).

In summary, the **contributions** of our work include:

- *The first real-time technique for spatial frequency adjustment of images and videos* (Chapter 3). It is based on a novel mathematical framework that enables its improved performance over the state of the art. It requires pre-processing the space-frequency decomposition and the unwrapped phase of the image, but frequency adjustment is subsequently performed in real-time.
- *A phase-unwrapping technique to recover the phases of real waves* that locally reconstruct an image (Chapter 4). Our solution naturally handles a sign ambiguity involving the phase unwrapping of real waves, in addition to simultaneously treating many unordered phase components without explicit sorting. As such, it extends the state-of-the-art on phase unwrapping that only covers a single phase component wrapped to $[-\pi, \pi)$.
- *A technique for processing color images for use with spatial frequency adjustment* (Section 4.3). It is based on a multi-channel optimization performed in RGB space and avoids some color artifacts associated with the use of PCA-based color reconstruction.

Figure 1.2 – Overview of our frequency-adjustment technique. (*Orange block*) The input image is decomposed into several cosine waves (Eq. (3.4)), represented by per-pixel scalar amplitudes and instantaneous-frequency vectors. The frequency vector direction and magnitude are given by the vector from the center of the color wheel to the pixel color in the wheel. This decomposition may be performed in the RGB space (applied to each color channel independently) or in some other color space where channels better encode the high-frequency content, such as the PCA-defined color space used for this illustration. (*Blue block*) The detected frequencies are then integrated by our method to recover a continuous scalar field of unwrapped phases. Note that the unwrapped phase values for each pixel are not sorted between the images, and as such the phase data may appear more discontinuous than they actually are. Note also the absence of phase information for low-detail regions, which are treated as residue only. (*Green block*) The frequency-adjusted content is obtained by reconstructing the waves from their original amplitudes and corresponding scaled unwrapped phases (Eq. (5.1)), using a user-provided scaling factor α (for this example, $\alpha = 0.25$). All waves are summed to obtain the frequency-adjusted content. Finally, to obtain the reconstructed image, the frequency-adjusted content is transformed back to RGB space (if initially in RGB space no transformation is needed) and summed back to the residual.

Source: The authors. “Barbara” image attributed to Allen Gersho (public domain).

2 RELATED WORKS AND TECHNICAL BACKGROUND

Our work was inspired by the *spectral remapping* (SR) technique for image downscaling (GASTAL; OLIVEIRA, 2017). SR remaps frequencies that are not representable in the target resolution to new representable frequencies, preserving the orientations of the original high-frequency patterns. This solution effectively encodes important structured details at lower frequencies during downscaling, instead of discarding them with an anti-aliasing filter. This results in more faithful representations of the original content. The technique, however, has only been demonstrated for image downscaling in combination with a remapping function that takes similarly-oriented waves to a common frequency. More importantly, (i) it is *non-interactive*, requiring the solution of a sparse linear system for each combination of input image I and downscaling factor r ; and (ii) the PCA-based reconstruction strategy used for processing color images may introduce some noticeable artifacts depending on the target resolution.

Our frequency-adjustment technique can be employed to image downscaling and addresses the aforementioned issues: it supports both decrease and increase of spatial frequencies, can be used with any image resizing factor $r > 0$, preserves the monotonic relation among the remapped frequencies, and can be used in interactive applications, being up to $100\times$ faster than SR. In addition, we also demonstrate its extension to videos.

Although our method can generate its results directly at the target resolution, it complements and can be used in combination with any resampling strategy, including downscaling and upscaling (KOPF; SHAMIR; PEERS, 2013; NEHAB; HOPPE, 2014; ÖZTIRELI; GROSS, 2015; SACHT; NEHAB, 2015; WEBER et al., 2016). Our technique can also be used with any single-image super-resolution method (e.g., (LIM et al., 2017; FREEDMAN; FATTAL, 2010; Yang et al., 2010)), and we demonstrate a hybrid solution combining super-resolution and spatial frequency adjustment in Chapter 6.

To achieve our results, we introduce a *phase unwrapping* algorithm that simultaneously handles several real-wave phase components, in addition to inter-component dependencies (Chapter 4). Phase unwrapping is the process of recovering a continuous function $\theta(x)$ from samples $\hat{\theta}(x)$ whose values are only known up to modulo 2π (OPPENHEIM; SCHAFER, 2010). This procedure is a fundamental building block in many applications, ranging from satellite topography and marine/earth seismology to the analysis of brain waves (CHILDERS; SKINNER; KEMERAIT, 1977; GOLDSTEIN; ZEBKER; WERNER, 1988). Phase unwrapping is tightly linked to the theory of instantaneous-

frequency decompositions and space/time-frequency analysis (MALLAT, 2009). In this context, it is often called “ridge” analysis, and it may be derived from Gabor (DELPART et al., 1992) and wavelet transforms (LILLY; OLHEDE, 2010). There is a significant amount of theory and algorithms developed in this area for 1-D signals, such as acoustic and speech data (GUILLEMAIN; KRONLAND-MARTINET, 1996). The corresponding 2-D/3-D problems (required for our frequency-adjustment technique) is significantly more challenging to solve (GHIGLIA; PRITT, 1998). Sections 2.2.3.1 and 2.2.3.2 include a more detailed discussion of this topic.

2.1 Image resizing

Image resizing is the process of representing an image in a different resolution while preserving the visual appearance of the original. It is usually separated in two cases, *image upscaling* and *image downscaling*. The first focuses in generating an image with a higher resolution than the original, and the second in generating a lower resolution image. Generally image upscaling produces good results with interpolation algorithms such as bilinear interpolation, however deep learning recently has been used for super-resolution applications (LIM et al., 2017). Image downscaling solutions traditionally involve the application of a low-pass filter to remove unrepresentable frequencies in the lower resolution, as these frequencies would introduce aliasing otherwise, and then apply a resampling algorithm to obtain a lower resolution image. Ideally, a sinc filter should be used to prevent aliasing, however, band-limited functions contain ringing artifacts that are distracting to the human eye. Furthermore, since it has infinite image-space support, several finite approximations have been used, such as the Lanczos kernel and several cubic filters (MITCHELL; NETRAVALI, 1988).

Modern works on image resizing use compact polynomial bases with recursive digital filters (NEHAB; HOPPE, 2014; SACHT; NEHAB, 2015) to obtain more efficient reconstruction algorithms. Alternative strategies try to preserve image details by locally optimizing bilateral kernels to the image content (KOPF; SHAMIR; PEERS, 2013). Other techniques model a maximization problem to increase the similarity between the image and its downsampled version (ÖZTIRELI; GROSS, 2015). More recently, an implementation of a downscaling algorithm for large images and videos was presented (WEBER et al., 2016), it was based on a variation of a joint-bilateral filter with a range kernel that favors color differences. We show examples of these downscaling algorithms in Figure 2.1.

Furthermore, like in image upscaling, deep learning models were also proposed for image downscaling, however the downsampled images do not preserve details. The idea of most deep learning models (LI et al., 2018; KIM et al., 2018; SUN; CHEN, 2020) is to generate a lower resolution image that when upsampled by super resolution deep learning methods generates a more similar image to the original. These methods usually work by training a resampling and a super resolution neural network together, by back-propagating the reconstruction error to both networks.

2.1.1 Overview on Spectral Remapping

The spectral remapping algorithm (GASTAL; OLIVEIRA, 2017) that inspired this work has four steps. First it detects high-frequency sinusoidal waves with the algorithm described in Section 2.2.2 (using a Short-Time Fourier Transform with a Gaussian window). Second the high-frequency waves have their frequencies remapped to a value close to the Nyquist limit of the downsampled resolution, but slightly lower to avoid ordinates that are distracting to the human eye and to avoid attenuating the frequencies. Thirdly, the phase-shifts are recomputed to align the waves as, after remapping the frequencies, the remapped waves may interfere incorrectly among themselves. This is done by minimizing the pointwise difference of neighboring waves (evaluated at a region located at the midpoint of the waves' centers), as a function of the desired remapped phases, weighted by a wave similarity measure. Finally, a new image is synthesized from the sum of the remapped waves and a residual (see Figure 5.1 (b)). The synthesis process consists in reconstructing the image from summing all remapped waves (windowed by the dual of the Gaussian window used for wave detection) and the residual to obtain an image with high-frequency content remapped to the desired frequency. The image is reconstructed in the original resolution as shown in Figure 2.1 ($\hat{s}$, top-right), thus a downscaling algorithm is necessary to generate the image in the desired resolution after remapping the frequencies. We compare our frequency-adjustment technique to Spectral Remapping in Chapter 6.

2.2 Technical Background

In this section we present an overview of the technical knowledge necessary to understand the concepts presented in this thesis.

Figure 2.1 – Several downscaling techniques applied to the image of Vittore at the top-left and the version with frequencies remapped by Spectral Remapping at the top-right. Resampling with: cubic O-MOMS (D), Öztireli and Gross (O), Lanczos (L), Lanczos with Spectral Remapping ($\dot{L}$), Weber et al. (W), Kopf et al. (K), Subsampling (III) and Subsampling with Spectral Remapping ($\ddot{\text{III}}$). Downscaling to 112×150 pixels.

Source: The authors.

Furthermore, we recommend the following reference material to fully understand the contents of this thesis:

- Oppenheim (OPPENHEIM; SCHAFER, 2010)
 - Chapter 2 "Discrete-Time Signals and Systems"
 - Chapter 8 "The Discrete Fourier Transform"
 - Chapter 10 "Fourier Analysis of Signals Using the Discrete Fourier Transform"
- Mallat (MALLAT, 2009)
 - Chapter 3 "Discrete Revolution"
 - Chapter 4 "Time Meets Frequency"
- Spectral Remapping (GASTAL; OLIVEIRA, 2017) in its entirety
- For a detailed review on the fundamental algorithms on phase unwrapping we also recommend (GHIGLIA; PRITT, 1998)

2.2.1 The discrete Fourier transform

Given a sequence of N complex numbers $x_n = x_0, \dots, x_{N-1}$, the discrete Fourier transform (DFT) maps x_n into another sequence $X_n = X_0, \dots, X_{N-1}$, defined by

$$X_k = \sum_{n=0}^{N-1} x_n f_k^*(n), \quad (2.1)$$

where

$$f_k(n) = \exp\left(\frac{2\pi k i}{N} n\right) \quad (2.2)$$

is a family of functions, where each f_k is a *complex exponential wave*, and f_k^* is the complex conjugate of f_k .

The sequence x_n may be recovered by the inverse discrete Fourier transform (IDFT) of X_n defined as

$$x_n = \frac{1}{N} \sum_{k=0}^{N-1} X_k f_k(n). \quad (2.3)$$

The absolute value of the scalar coefficient multiplying a sinusoidal wave is called the *amplitude* of the wave, the argument of the wave is called the *phase* and the derivative

of the phase is called the *instantaneous frequency* of the wave. Thus the complex wave

$$f_k(n) = \exp\left(\frac{2\pi k i}{N}n\right) = \left[\cos\left(\frac{2\pi k}{N}n\right) + i \sin\left(\frac{2\pi k}{N}n\right)\right]$$

has amplitude $A = 1$, phase $\theta(n) = \frac{2\pi k}{N}n$ and instantaneous frequency $\frac{d\theta(n)}{dn} = \frac{2\pi k}{N}$.

2.2.2 Space-frequency decomposition: Wave detection

Our technique begins by obtaining a non-harmonic space-frequency decomposition. For this, a wave detection algorithm based on the Gabor transform, a short time Fourier transform that uses a Gaussian window, is used to determine local frequency content of regions of a signal. The transform works by applying a sliding Gaussian window to the signal to isolate the localized sinusoidal waves in each region of the image. For each windowed region, a Fourier transform detects all waves present in the region (see Figure 2.2).

The wave detection algorithm is described by Gastal and Oliveira (GASTAL; OLIVEIRA, 2017), it consists in taking a set of discrete samples from the windowed image, computing its discrete Fourier Transform (DFT), finding the indices (k, l) of the maximum absolute coefficient of the DFT and fitting a quadratic surface around a 3×3 neighborhood of (k, l) . The horizontal and vertical frequencies are then given by the indices (a, b) of the surface's peak (in frequency units) and the amplitude A by the height of the peak. Then, the phase-shift c of the wave is found by solving a minimization problem of the difference from the wave and the windowed image. The wave is then fully representable in the form $G(x, y)A \cos(ax + by + c)$, where G is the Gaussian window used to detect the wave.

After detecting a wave, it is removed by subtracting its values from the image and the process is repeated until either: a certain number of waves are detected (in our implementation we use 10 waves), the amplitude of the last detected wave is very small or the detected wave is not reliable (i.e., because the quadratic does not have a maximum, has no maximum in the 3×3 neighborhood or the decomposition is highly non-orthogonal and badly conditioned). We note that the number of components detectable is at most half the size of the detection window (half of the components in the window are the complex-conjugate of the others), however such a case is unlikely as the image would need to be extremely noisy to contain so many components.

Figure 2.2 – Example of windowing used for frequency detection. The tablecloth in (a) is transformed to a PCA-defined space and sampled at the center of the zoomed tablecloth in (b). Then a (truncated) Gaussian window is multiplied in (c) to isolate the localized high-frequency content as shown in the Fourier amplitude spectrum (d).

Source: (a) “Barbara” image attributed to Allen Gersho (public domain). (b–d) The authors.

For the purpose of obtaining a non-harmonic space-frequency decomposition of the image h , we employ this wave detection algorithm to separate the image as $h = s + d$, where s contains most of the high-frequency information of h , given by summing the detected waves that exceed a certain predefined frequency threshold (Figure 5.1), and d is the residual from subtracting s from h (GASTAL; OLIVEIRA, 2017).

2.2.3 Phase Unwrapping

Consider a continuous complex wave signal $\exp(i\theta(x, y))$ described by a continuously-differentiable phase function θ . Phase unwrapping is the process of recovering the continuous θ from a phase function $\hat{\theta}$ that is *wrapped* by 2π (that is, $\hat{\theta} = \theta \bmod 2\pi$). This occurs because $\hat{\theta}$ is obtained as the argument of the periodic function, $\hat{\theta} = \arg(\exp(i\theta))$, and thus its value is in a single period interval. As such, $\hat{\theta}$ is related to the true phase θ by the expression

$$\theta(x, y) = \hat{\theta}(x, y) + 2\pi k(x, y), \quad (2.4)$$

for unknown $k(x, y) \in \mathbb{Z}$.

By noting that the gradient fields of $\nabla\theta$ and $\nabla\hat{\theta}$ are equal ($\nabla\theta = \nabla\hat{\theta}$), phase-unwrapping amounts to computing the vector-field integral “ $\int \nabla\hat{\theta}$ ” to obtain the scalar potential field $\tilde{\theta} = \theta + C$, for some integration constant C , where $\tilde{\theta}$ is smooth.

2.2.3.1 Unwrapping the phase of one complex wave

The phase-unwrapping literature deals with wrapped phase values $\hat{\theta}(x, y) \in [-\pi, \pi)$ that are related to the true phase θ according to Eq. (2.4). In practice, this is equivalent to a phase obtained from a *complex* wave $\exp(i\theta(x, y))$. Note that the wrapped phase recovered from a complex wave is given by $\hat{\theta} = \arg(\exp(i\theta)) \in [-\pi, \pi)$ (ITOH, 1982). Since the integer shifts $k(x, y)$ in Eq. (2.4) define a piecewise-constant function, the unwrapped and wrapped gradients are such that

$$\nabla\theta(x, y) = \nabla\hat{\theta}(x, y) \quad \text{almost everywhere.} \quad (2.5)$$

If the vector field $\nabla\hat{\theta}(x, y)$ is conservative, then the scalar potential field $\theta(x, y)$ is directly determined by the path integral

$$\theta(x, y) = \int_{t_0}^{t_1} \nabla\theta(\mathbf{v}(t)) \cdot \mathbf{v}'(t) dt \quad (2.6)$$

for some path $\mathbf{v} : [t_0, t_1] \rightarrow \mathbb{R}^2$ with $\mathbf{v}(t_0) = \mathbf{v}_0$ and $\mathbf{v}(t_1) = (x, y)$ (up to some integration constant) (TAKAJO; TAKAHASHI, 1988).

In the discrete setting, the wrapped phase $\hat{\theta}(x, y)$ is only sampled at integer pixel locations, and the discrete analogue to Eq. (2.5) is given by finite differences. Some extra care must be taken, however, to handle the discontinuities caused by the integer shifts $k(x, y)$. Section 2.2.3.2 reviews how this is done in the phase-unwrapping literature (ITOH, 1982) (note however that in our method we use an alternative approach to obtain estimates for the phase’s gradient, discussed in Section 2.2.4). The end result is an estimated discrete gradient field $\nabla\tilde{\theta}$ that approximates the true field $\nabla\theta$ at each pixel (x, y) , and the unwrapped phase $\theta(x, y)$ is obtained by integrating $\nabla\tilde{\theta}$. While, in theory, this may be performed by a cumulative sum over any path covering all pixels in the image, this process is sensitive to noise and errors in the measurements (ZEBKER; LU, 1998), in addition to ambiguities inherent to the phase-unwrapping problem (SCIVIER; FIDDY, 1985). Thus, in practice, the integration process becomes path dependent and a good choice

of integration path must be made. Modern approaches obtain good results by the use of combinatorial optimization with carefully-selected heuristics (HERSZTERG; POGGI; VIDAL, 2019). Unfortunately, the time complexity of these methods scale as $O(N^3)$ in the number of pixels.

A second class of algorithms work by integrating the field $\nabla\tilde{\theta}$ in the least-squares sense (TAKAJO; TAKAHASHI, 1988; SONG et al., 1995; ZEBKER; LU, 1998). This is equivalent to solving a discrete Poisson problem. The benefit of these approaches is that they recover smoothly-varying phase information even in regions of low signal-to-noise ratio, but at the cost of some distortion in the recovered phase (which may not exactly match the wrapped measurements) (ZEBKER; LU, 1998). For our frequency-adjustment technique, obtaining a smooth phase field is of utmost importance, as any discontinuity tends to introduce objectionable visual artifacts. As such, we opt for a least-squares solution, discussed in Sections 4.1.1 and 4.2. Since this amounts to solving a linear system, the unwrapping computation takes at most a few seconds for a typical image at full HD resolution on commodity hardware. The actual time depends on the image resolution and its frequency content. Note, however, that recovering $\theta(x, y)$ can be done as pre-processing, after which one can perform spatial frequency adjustments in real-time using independent per-pixel operations (Chapter 5).

2.2.3.2 Common approach for estimating the wrapped gradient field for one complex wave

The following relation holds between the phase θ of a complex wave and its wrapped version $\hat{\theta}(x, y) = \arg(\exp(i\theta(x, y)))$ (ITOH, 1982):

$$\nabla\theta(x, y) \approx \text{wrap}(\Delta\hat{\theta}(x, y)) \quad \text{if} \quad \nabla\theta(x, y) \in [-\pi, \pi]^2; \quad (2.7)$$

where $\Delta = (\Delta_x, \Delta_y)$ gives the discrete gradient, for

$$\begin{aligned} \Delta_x\hat{\theta}(x, y) &= \hat{\theta}(x, y) - \hat{\theta}(x - 1, y), \text{ and} \\ \Delta_y\hat{\theta}(x, y) &= \hat{\theta}(x, y) - \hat{\theta}(x, y - 1). \end{aligned}$$

The operator $\text{wrap}(\theta) = \theta + 2\pi \lfloor (\pi - \theta)/2\pi \rfloor$ operates elementwise when applied to a vector (HERSZTERG; POGGI; VIDAL, 2019). The identity in Eq. (2.7) holds everywhere if the horizontal and vertical variations in θ are bounded by π radians (ITOH, 1982). This restriction, that $\nabla\theta(x, y) \in [-\pi, \pi]^2$, is trivially satisfied for sampled complex signals. The continuous signal reconstruction with sinc interpolation is naturally band-limited to the

Nyquist frequency interval $[-\pi, \pi)^2$. This implies that the components of the instantaneous frequency $\nabla\theta(x, y)$ are in $[-\pi, \pi)$. As such, given measurements $\hat{\theta}(x, y)$ of the wrapped phase, the unwrapped phase $\theta(x, y)$ for one complex wave can be obtained by integrating $\text{wrap}(\Delta\hat{\theta}(x, y))$ (ITOH, 1982; SONG et al., 1995).

2.2.4 Estimating the phase gradient field for one real wave

Our frequency-adjustment method requires computing the gradient values $\nabla\theta(x, y)$. This section shows how these can be obtained (up to an unknown sign) from the decomposition of Gastal and Oliveira (GASTAL; OLIVEIRA, 2017).

The space-frequency decomposition presented in Section 2.2.2 approximates the (Gaussian-windowed) neighborhood around each pixel as a collection of plane waves of the form $A \cos(ax + by + c)$, where $(a, b) \in [-\pi, \pi)^2$ are constant horizontal and vertical frequencies, $c \in [0, 2\pi)$ is a constant phase-shift, and $A > 0$ is a constant amplitude. In other words, assuming a single phase component (the case of several components is discussed in Section 4.2), a tuple of measurements (a, b, c) is computed for each pixel (x_0, y_0) , and the approximation

$$s(x_0 + \delta_x, y_0 + \delta_y) \approx A \cos(\hat{\omega}) = A \cos(a\delta_x + b\delta_y + c) \quad (2.8)$$

is a linear-phase approximation of the image's high-frequency content s around pixel (x_0, y_0) . The (unknown) true phase value at (x_0, y_0) matches the measured phase-shift c up to modulo 2π (i.e., c is a wrapped measurement).

By definition, a function and its best linear approximation at point (x_0, y_0) have matching first derivative at such point. Thus, for the desired representation $s(x, y) = A \cos(\theta(x, y))$ (presented in Chapter 3, Eq. (3.2)), with a possibly non-linear phase function $\theta(x, y)$, and given Eq. (2.8), with (a, b) computed for a particular pixel (x_0, y_0) , one has

$$\nabla\theta(x_0, y_0) = \nabla\hat{\omega} = (a, b). \quad (2.9)$$

Moreover, since $\cos(\hat{\omega}) = \cos(-\hat{\omega})$, it is also true that

$$s(x + \delta_x, y + \delta_y) \approx A \cos(-\hat{\omega}) = A \cos(-a\delta_x - b\delta_y - c); \quad (2.10)$$

so it could also be the case that we have a sign flip:

$$\nabla\theta(x_0, y_0) = -\nabla\hat{\omega} = (-a, -b). \quad (2.11)$$

In other words, $\nabla\theta(x_0, y_0) = \sigma(x_0, y_0) \nabla\hat{\varphi}(x_0, y_0)$ for an unknown sign $\sigma(x_0, y_0) \in \{-1, +1\}$, where the measured phase gradient

$$\nabla\hat{\varphi}(x_0, y_0) = (a, b) \quad (2.12)$$

is given by the horizontal and vertical plane-wave frequencies a and b , detected at pixel (x_0, y_0) , by a non-harmonic Gabor space-frequency analysis (Section 2.2.2).

3 FREQUENCY ADJUSTMENT

Consider a continuous wave signal $s(x, y) = \cos(\theta(x, y))$ described by a continuously-differentiable phase function θ . The horizontal $\nabla_x \theta(x, y)$ and vertical $\nabla_y \theta(x, y)$ *instantaneous frequencies* of this signal are given by the partial derivatives of its phase (MALLAT, 2009):

$$\nabla_x \theta(x, y) = \frac{\partial}{\partial x} \theta(x, y) \quad \text{and} \quad \nabla_y \theta(x, y) = \frac{\partial}{\partial y} \theta(x, y). \quad (3.1)$$

Similarly, the local wavefront orientation (normal direction) at (x, y) is given by the gradient vector $\nabla \theta(x, y) = [\nabla_x \theta(x, y), \nabla_y \theta(x, y)]$.

A *frequency-adjusted* version s_α of s , where frequencies are scaled by a factor $\alpha > 0$, is obtained as:

$$s_\alpha(x, y) = \cos(\alpha \theta(x, y)). \quad (3.2)$$

It is easy to see that the horizontal and vertical frequencies of s_α are $\alpha \nabla_x \theta$ and $\alpha \nabla_y \theta$, i.e., scaled versions of the corresponding frequencies of s . The local direction of wave propagation is preserved in s_α since $\nabla \alpha \theta = \alpha \nabla \theta$. Finally, s_α is still a continuous function, as was s . Figure 3.1 illustrates an example of frequency adjustment for a signal where θ is known analytically.

3.1 Frequency Adjustment of Natural Images

Given a natural image $h(x, y)$, we seek a decomposition which represents it as a sum of N components with phase functions $\theta^1 \dots \theta^N$, amplitude functions $A^1 \dots A^N$, and a residual $d(x, y)$:

$$h = s + d \quad \text{where} \quad s(x, y) = \sum_{n=1}^N A^n(x, y) \cos(\theta^n(x, y)). \quad (3.3)$$

Obtaining such a decomposition for a given image h is a severely under-constrained problem. Thus, we require the summation in $s(x, y)$ to only encode high-frequency content that should be affected by the frequency adjustment (remaining content goes into the residual d). Furthermore, the phase functions are required to be piecewise smooth for the instantaneous frequencies to be well defined almost everywhere, making $s(x, y)$ (and thus h) amenable to frequency adjustment. Finally, the amplitudes $A^n(x, y)$ should vary slowly when compared to local frequencies, so that high-frequency variations become

Figure 3.1 – Example of frequency adjustment. The continuous phase function $\theta(x, y)$ is defined by a Perlin noise scalar field (left). Scaling the phase by factors $\alpha > 0$ results in frequency-adjusted signals $\cos(\alpha\theta(x, y))$, which preserve the local wavefront orientations.

Source: The authors.

encoded in the phases only (MALLAT, 2009).

The decomposition shown in Eq. (3.3) is related to the Hilbert-Huang transform (HHT), that separates a signal into a collection of so-called Intrinsic Mode Functions (IMFs) and a residue (HUANG et al., 1998; HUANG, 2014). The HHT differs from our decomposition by the way the phases are computed: while we use windowed Fourier transforms (Gabor analysis) to obtain instantaneous frequencies (GASTAL; OLIVEIRA, 2017) (which are integrated to generate the unwrapped phase), HHT uses the Hilbert transform. Both decompositions however require unwrapping to obtain a continuous phase, which we discuss in the following sections.

For a frequency adjustment factor $\alpha > 0$, the frequency-adjusted image h_α is obtained by using the scaled phase functions $\alpha\theta^n$ during reconstruction, as in Eq. (3.2). Figure 1.2 presents an overview of the full frequency-adjustment processing pipeline.

3.2 Finding a Decomposition

Time/Space-frequency analysis may be employed to measure the instantaneous frequencies of several non-harmonic spectral components (MALLAT, 2009). Gastal and Oliveira (GASTAL; OLIVEIRA, 2017) describe a practical algorithm for this purpose, approximating the neighborhood around each pixel as a sum of Gaussian-windowed plane waves. From the constant frequencies and phase-shifts of such waves it is possible to extract (Section 2.2.4), for each pixel, a series of phase values $\hat{\varphi}^n(x, y)$ and associated amplitudes $B^n(x, y)$, $n = 1 \dots N$, such that

$$h = s + d \quad \text{where} \quad s(x, y) = \sum_{n=1}^N B^n(x, y) \cos(\hat{\varphi}^n(x, y)). \quad (3.4)$$

Figure 3.2 – Example of unsuccessful frequency adjustment. When the phase function $\hat{\varphi}(x, y)$ is wrapped (values recovered as $\hat{\varphi} = \arccos(\cos(\varphi))$ are in $[0, \pi]$ – see color scale on the left), scaling the phase by factors $\alpha > 0$ does **not** result in frequency-adjusted signals. One must therefore **unwrap** the phase function before being able to perform frequency adjustment. Compare to Figure 3.1.

Source: The authors.

Despite the similarity of this expression to that of Eq. (3.3), the values $\hat{\varphi}^n$ cannot be directly employed for frequency adjustment since they are *wrapped* to $[0, \pi]$. The problem that occurs when trying to do so is illustrated in Figure 3.2.

The occurrence of wrapped values is an inherent mathematical limitation of phase measurements. Since $\cos(\theta) = \cos(\theta + 2\pi m)$ for $m \in \mathbb{Z}$, the phase θ is not uniquely defined. Moreover, when trying to recover θ as $\arccos(\cos(\theta))$, one obtains the principal value $\hat{\varphi} \in [0, \pi]$. Recovering a continuous and differentiable phase function $\theta(x, y)$ from wrapped values $\hat{\varphi}(x, y)$ is called *phase unwrapping* (Section 2.2.3) (OPPENHEIM; SCHAFER, 2010). The measured phase values $\hat{\varphi}^n$ recovered from space-frequency analysis techniques (including Gastal and Oliveira's) are wrapped, and therefore require unwrapping in order to define phase functions θ^n amenable to frequency adjustment.

4 PHASE UNWRAPPING IN TWO DIMENSIONS

While 1-D phase unwrapping has a simple solution (MERTZ, 1979), 2-D unwrapping is more involving and requires specialized algorithms (GHIGLIA; PRITT, 1998). It has been extensively studied due to its relevance to many fields, including interferometry and magnetic resonance imaging, and several algorithms exist aimed at these and other applications (we include a detailed review in Section 2.2.3). *These algorithms work with a wrapped phase $\hat{\theta} \in [-\pi, \pi)$.* Thus, solving the phase unwrapping problem in this case consists of finding the shift factors $k(x, y)$ in Eq. (2.4).

Digital images are real functions and modeled using real waves: the cosines in Eq. (3.3). Therefore, in this situation the measured phase *has a reduced angular range*, $\hat{\varphi} = \arccos(\cos(\theta)) \in [0, \pi]$, and an *ambiguous sign*, since $\cos(\hat{\varphi}) = \cos(-\hat{\varphi})$. As a result, $\hat{\varphi}$ is related to the true phase θ by an unknown sign, σ , in addition to an unknown shift, k (the diagram to the right illustrates the case $k = 2, \sigma = +1$):

$$\theta(x, y) = \sigma(x, y) \hat{\varphi}(x, y) + 2\pi k(x, y) \text{ for } \sigma(x, y) \in \{-1, +1\}. \quad (4.1)$$

Such *sign ambiguity* and *reduced angular range* make phase unwrapping for digital images and videos a challenging task.

In the following sections we present our novel 2-D phase unwrapping algorithm designed for spatial frequency adjustment of digital images (its extension to videos, which requires 3-D processing, is described in Section 5.3). Our solution handles the following issues that arise when dealing with real waves: (i) phase measurements with ambiguous sign, due to the conjugate-symmetry of the spectrum of real signals (Section 4.1); (ii) presence of several unordered phase components, required for representing complex patterns (Section 4.2); and (iii) correlated multichannel components in color images (Section 4.3). Furthermore, implementation details and optimizations are described in Chapter 5.

4.1 Unwrapping the phase of one real wave

For frequency adjustment of digital images, one is interested in unwrapping the phase $\theta(x, y)$ from a real wave component $\cos(\theta(x, y))$. The wrapped phase measurement

$\hat{\varphi}(x, y) \in [0, \pi]$ is thus related to the true phase according to Eq. (4.1). Since $k(x, y)$ defines a piecewise-constant function, the unwrapped and wrapped gradients are related only by the unknown sign $\sigma(x, y)$ at each position (x, y) :

$$\nabla\theta(x, y) = \sigma(x, y)\nabla\hat{\varphi}(x, y) \quad \text{almost everywhere.} \quad (4.2)$$

In Section 2.2.4 we showed how to obtain $\nabla\hat{\varphi}$ directly from a space-frequency decomposition. We now focus on how to integrate the gradient field $\nabla\theta = \sigma\nabla\hat{\varphi}$ to recover the phase function, *without explicitly computing the signs* σ .

(Throughout the text we denote the pointwise product of two functions f and g as fg ; that is, $(fg)(x, y) = f(x, y)g(x, y)$. Furthermore, $f = g$ denotes pointwise equality: $f(x, y) = g(x, y), \forall(x, y)$.)

4.1.1 Integrating the phase gradient field for one real wave.

To recover the unwrapped phase θ , the gradient field $\nabla\theta = \sigma\nabla\hat{\varphi}$ is integrated in the least-squares sense. Although the signs $\sigma(x, y)$ are unknown, our insight is that it is possible to check whether neighboring pixels experience a change in sign (Section 4.1.2). This information is sufficient for integrating $\nabla\theta$ in a sign-consistent way through some clever algebraic manipulations, shown below.

Let T_x and T_y be two unit-shift operators that act on a function θ as $T_x(\theta)(x, y) = \theta(x - 1, y)$, and $T_y(\theta)(x, y) = \theta(x, y - 1)$. Furthermore, let $\circ$ in $T_\circ$ stand for either x or y . Then, a trapezoidal integration rule gives the following linear relations between the unknown phase values $\theta(x, y)$ of all pixels (SOLOMON, 2015):

$$\theta - T_\circ(\theta) = \frac{1}{2} [\nabla_\circ\theta + T_\circ(\nabla_\circ\theta)]. \quad (4.3)$$

Given that $\nabla\theta = \sigma\nabla\hat{\varphi}$, Eq. (4.3) may be expanded as (we use underlines to highlight the parts of the equations that have changed):

$$\theta - T_\circ(\theta) = \frac{1}{2} [\underline{\sigma\nabla_\circ\hat{\varphi}} + T_\circ(\underline{\sigma\nabla_\circ\hat{\varphi}})]. \quad (4.4)$$

By introducing a function φ satisfying $\sigma\varphi \stackrel{\text{def}}{=} \theta$, Eq. (4.4) becomes

$$\underline{\sigma\varphi} - T_\circ(\underline{\sigma\varphi}) = \frac{1}{2} [\sigma\nabla_\circ\hat{\varphi} + T_\circ(\sigma\nabla_\circ\hat{\varphi})]. \quad (4.5)$$

Note that to solve the above equation one must know σ . To remove this dependency, we start by multiplying both sides by σ :

$$\underline{\sigma}^2 \varphi - \underline{\sigma} T_{\circ}(\sigma \varphi) = \frac{1}{2} [\underline{\sigma}^2 \nabla_{\circ} \hat{\varphi} + \underline{\sigma} T_{\circ}(\sigma \nabla_{\circ} \hat{\varphi})]. \quad (4.6)$$

Note that $\sigma^2 = 1$ for any (x, y) , since $\sigma(x, y) \in \{-1, +1\}$. Furthermore, since $T_{\circ}(fg) = T_{\circ}(f)T_{\circ}(g)$,^{*} Eq. (4.6) may be simplified to

$$\varphi - \underline{\Delta_{\circ} \sigma} T_{\circ}(\varphi) = \frac{1}{2} [\nabla_{\circ} \hat{\varphi} + \underline{\Delta_{\circ} \sigma} T_{\circ}(\nabla_{\circ} \hat{\varphi})], \quad (4.7)$$

where $\Delta_{\circ} \sigma \stackrel{\text{def}}{=} \sigma T_{\circ}(\sigma)$. Note that $\Delta_x \sigma$ and $\Delta_y \sigma$ encode the locations where σ changes sign. More precisely, $\Delta_x \sigma(x, y) = -1$ whenever the sign of $\sigma(x, y)$ differs from the sign of the neighboring pixel $\sigma(x - 1, y)$, and $\Delta_x \sigma(x, y) = +1$ otherwise.[†] An analogous relation is true for $\Delta_y \sigma(x, y)$ in the y -dimension.

As shown below in Section 4.1.2, it is possible to compute $\Delta_{\circ} \sigma$ despite σ being unknown. As such, the only unknowns in Eq. (4.7) are the values of the function φ , which may be found by solving the associated linear system through least squares:

$$\min_{\varphi} \sum_{\circ \in \{x, y\}} \|\varphi - \Delta_{\circ} \sigma T_{\circ}(\varphi) - f_{\circ}\|^2 + \lambda \|\varphi\|^2, \quad (4.8)$$

where $f_{\circ}(x, y)$ encodes the right-hand-side of Eq. (4.7), for $\circ \in \{x, y\}$. The recovered phase φ is enough to perform frequency adjustment, without knowing the per-pixel sign σ , since the cosine is an even function. Thus $\cos(\alpha \varphi) = \cos(\alpha \sigma \varphi) = \cos(\alpha \theta)$ and so it is not necessary to recover θ .

The minimum-norm regularization term $\|\varphi\|^2$ in Eq. (4.8) is required for two reasons. First, without regularization, the existence of one solution φ would imply an infinite number of solutions $\varphi + C$, for any real integration constant C . Second, due to the inherent trade-off between space and frequency locality (MALLAT, 2009), the numerical space-frequency analysis only detects phase information for high-frequency portions of the image (GASTAL; OLIVEIRA, 2017). As a result, the measurements $\nabla \hat{\varphi}$ and consequently $f_{\circ}$ have undefined values for some pixels, which may result in unconstrained unknowns. The same occurs along the boundary of the image domain. The regularizer forces any

^{*} In 1-D: $T(fg)(x) = (fg)(x - 1) = f(x - 1)g(x - 1) = T(f)(x)T(g)(x)$.

[†] From the definition of the unit shifts T_x and T_y :

$$\begin{aligned} \Delta_x \sigma(x, y) &= \sigma(x, y) T_x(\sigma)(x, y) = \sigma(x, y) \sigma(x - 1, y), \text{ and} \\ \Delta_y \sigma(x, y) &= \sigma(x, y) T_y(\sigma)(x, y) = \sigma(x, y) \sigma(x, y - 1). \end{aligned}$$

unconstrained variables to zero. Finally, an additional benefit of the regularizer is that it penalizes the occurrence of large absolute values in φ , making the frequency-adjusted reconstruction $\cos(\alpha\varphi)$ less sensitive to small changes in α (the change in $\cos(\alpha\varphi)$ when varying α is given by $\frac{\partial}{\partial\alpha} \cos(\alpha\varphi) = -\varphi \sin(\alpha\varphi) \propto \varphi$). This is important in interactive applications where α is user-controlled.

Eq. (4.8) has a unique solution obtainable from the linear system

$$(D_x^T D_x + D_y^T D_y + \lambda I) \varphi = D_x^T f_x + D_y^T f_y, \quad (4.9)$$

where I is the identity matrix, and D_x and D_y are similar to backward-difference matrices, but (i) have some entries with flipped signs (according to $\Delta_x \sigma$ and $\Delta_y \sigma$); and (ii) only include constraints for non-missing values (the reconstruction of pixels with missing values is handled by windowing, discussed in Chapter 5). We use $\lambda = 10^{-6}$ in the regularizer since it produces good results according to our empirical tests.

4.1.2 Determining sign changes and computing $\Delta\sigma$.

Let p and q be two neighboring pixels, for which we compute local gradients $\nabla \hat{\varphi}(p)$ and $\nabla \hat{\varphi}(q)$ according to Section 2.2.4. Also, recall that $\nabla \theta = \sigma \nabla \hat{\varphi}$. Two situations are possible:

1. **There is** a sign change between the phase gradients $\nabla \theta(p)$ and $\nabla \theta(q)$ of p and q , meaning that $\sigma(p) \neq \sigma(q)$;
2. **There is no** sign change between p and q , and $\sigma(p) = \sigma(q)$.

Possibility (1) implies $\Delta\sigma(p) = \sigma(p) \sigma(q) = -1$, while possibility (2) implies $\Delta\sigma(p) = \sigma(p) \sigma(q) = +1$. To determine which one is true, we recall that the phase θ should be continuously differentiable, meaning that the phase gradient $\nabla \theta$ should be continuous. As such, we choose the option that better preserves the continuity of $\nabla \theta$, i.e., that results in the smallest change between $\nabla \theta(p)$ and $\nabla \theta(q)$. Thus,

$$\Delta\sigma(p) = \arg \min_{v \in \{-1, +1\}} \|\nabla \hat{\varphi}(p) - v \nabla \hat{\varphi}(q)\|^2. \quad (4.10)$$

4.2 Unwrapping the phases of several real waves

Natural images contain complex structured patterns, such as in the honeycomb and sunflower photographs of Figure 6.1. These patterns require the summation of several cosine components in order to be correctly represented by the decomposition of Eq. (3.3). The several phase functions $\{\theta^1, \theta^2, \dots\}$ that build-up a pattern are in principle independent and could be separately unwrapped. In practice, however, this independent unwrapping is severely error prone. To understand why, note that, before unwrapping, one would be required to assign each n -th gradient measurement $\nabla \hat{\varphi}^n(p)$, of a pixel p , to a particular phase component $\theta^{\kappa_p(n)}$, where $\kappa_p(n)$ is an indexing function (permutation) specific for p . Furthermore, the permutations κ_p and κ_q for all neighboring pixels p and q should be consistent to each other, assigning the gradients that belong together to the same phase component. Otherwise, severe unwrapping errors will occur.

Due to ambiguous information returned by the space-frequency decomposition, sometimes it is not possible to precisely define the permutations κ_p for all pixels p . Furthermore, for particular pixels, the permutations are often incorrectly defined due to noise or frequency measurement errors. As a result, the whole phase unwrapping solution becomes compromised.

We avoid the aforementioned difficulties by, instead, defining weighted associations between *all* measurements $\nabla \hat{\varphi}^n$ and *all* phase functions θ^n (in place of discrete permutations). This significantly reduces the possibility of unwrapping errors, especially in the presence of ambiguous information returned by the space-frequency decomposition. It also makes the algorithm more robust to noise and measurement errors, and handles missing phase information.

Consider a pixel p and let $N(p)$ denote the number of measured phase components for p . Furthermore, let $T_o(p)$ denote the neighboring pixel of p according to the unit-shift operator T_o .[‡] For valid indices $n = 1 \dots N(p)$ and $m = 1 \dots N(T_o(p))$, the inter-component version of the integration rule (Eq. (4.7)), which couples the n -th component of pixel p to the m -th component of the neighboring pixel $T_o(p)$, is given by

$$\varphi^n - \Delta_o^{nm} \sigma T_o(\varphi^m) = \frac{1}{2} \left[\nabla_o \hat{\varphi}^n + \Delta_o^{nm} \sigma T_o(\nabla_o \hat{\varphi}^m) \right]. \quad (4.11)$$

[‡]Consider T_o as an overloaded shift operator that may receive either a function or a pixel $p = (x, y)$. Thus $T_x(p) = T_x((x, y)) = (x-1, y)$ and $T_y(p) = T_y((x, y)) = (x, y-1)$. If one is using sparse sampling, as discussed in Section 5.1, the shift operators will perform shifts of τ units.

As before, $\Delta^{nm}\sigma(p)$ determines if there is an expected change in sign between the n -th gradient measurement $\nabla\hat{\varphi}^n(p)$ of p , and the m -th measurement $\nabla\hat{\varphi}^m(q)$ of $q = T_\circ(p)$.

Some components are more likely to belong together than others. Thus, let the weight μ_{pq}^{nm} (defined below) be proportional to the probability of $\nabla\hat{\varphi}^n(p)$ and $\nabla\hat{\varphi}^m(q)$ belonging to the same phase component θ^k . We transform Eq. (4.11) into a *weighted* least-squares system, that simultaneously searches for the unwrapped phase components φ^n that best adapt to *all* couplings between measurements, but giving greater weight to the most probable ones (through μ^{nm}):

$$\min_{\varphi^n} \sum_{\circ \in \{x,y\}} \sum_{\forall n} \sum_{\forall m} \left\| \mu_{\circ}^{nm} \{L_{\circ}^{nm} - f_{\circ}^{nm}\} \right\|^2 + \lambda \sum_{\forall n} \|\varphi^n\|^2. \quad (4.12)$$

Here, $L_{\circ}^{nm}$ is the left-hand-side of Eq. (4.11), and $f_{\circ}^{nm}$ is its right-hand side. Furthermore, it is understood that n and m only vary over the valid indices for each pixel. For μ_{pq}^{nm} , we use the wave alignment measure proposed by Gastal and Oliveira (GASTAL; OLIVEIRA, 2017) (which determines how well the local Gabor approximations, associated with $\nabla\hat{\varphi}^n(p)$ and $\nabla\hat{\varphi}^m(q)$, match around the midpoint between p and q), scaled by a frequency-vector orientation similarity measure (Eq. (4.17)). That is

$$\mu_{pq}^{nm} = \text{Ang}(\nabla\hat{\varphi}^n(p), \nabla\hat{\varphi}^m(q)) \exp\left(-\frac{\delta^2}{\lambda^2}\right), \quad (4.13)$$

where $\exp\left(-\frac{\delta^2}{\lambda^2}\right)$ is a Gaussian function with variance $\lambda = 0.5$ (GASTAL; OLIVEIRA, 2017) and

$$\delta = \frac{\sqrt{\sum_{v \in \Upsilon} |B^n(p) \cos(\nabla\hat{\varphi}^n(p) \cdot v + c^n(p)) - B^m(q) \cos(\nabla\hat{\varphi}^m(q) \cdot v + c^m(q))|^2}}{\min(B^n(p), B^m(q))}$$

where Υ is the 3×3 neighborhood around $\frac{p+q}{2}$ and $c^n(p)$ is the phase-shift of the n -th wave at p (see Sections 2.2.2 and 2.2.4).

4.3 Processing Color Images

Frequency adjustment should be performed in all channels of a color image, since they have correlated information (REINHARD et al., 2001). Thus, while each color channel

has its own phase decomposition (Eq. (3.3)), their phase functions are unwrapped simultaneously, in a single optimization step. Let $\mathcal{E}_c$ be the energy functional from Eq. (4.12), written for the phase components φ_c^n of a specific color channel $c \in \{R, G, B\}$. We minimize the global energy

$$(\mathcal{E}_R + \mathcal{E}_G + \mathcal{E}_B) + \beta (\mathcal{X}_{RG} + \mathcal{X}_{GB} + \mathcal{X}_{RB}), \quad (4.14)$$

where $\mathcal{X}_{c_1 c_2}$ are cross-channel phase constraints, between two channels c_1 and c_2 , and $\beta = 10^{-3}$ is an empirically chosen constant that generates good results. These new constraints seek similar phase solutions $\varphi_{c_1}^n(p)$ and $\varphi_{c_2}^m(p)$, for a pixel p , if the measured instantaneous frequencies $\nabla \hat{\varphi}_{c_1}^n(p)$ and $\nabla \hat{\varphi}_{c_2}^m(p)$ are similar, between channels c_1 and c_2 . The indices $n = 1 \dots N_{c_1}(p)$ and $m = 1 \dots N_{c_2}(p)$ range over all detected phase components for pixel p at, respectively, color channels c_1 and c_2 .

The cross-channel constraints $\mathcal{X}_{c_1 c_2}$ are expressed as

$$\mathcal{X}_{c_1 c_2} = \sum_{n=1}^{N_{c_1}} \sum_{m=1}^{N_{c_2}} \eta_{c_1 c_2}^{nm} (\varphi_{c_1}^n - \Delta_{c_1 c_2}^{nm} \sigma \varphi_{c_2}^m)^2, \quad (4.15)$$

where $\Delta_{c_1 c_2}^{nm} \sigma(p)$ determines if there is an expected change in sign between the gradient vectors $\nabla \hat{\varphi}_{c_1}^n(p)$ and $\nabla \hat{\varphi}_{c_2}^m(p)$. Similarly, the weight $\eta_{c_1 c_2}^{nm}(p)$ measures the similarity of such vectors, according to their orientations and magnitudes:

$$\eta_{c_1 c_2}^{nm} = \text{Ang}(\nabla \hat{\varphi}_{c_1}^n, \nabla \hat{\varphi}_{c_2}^m) \text{Mag}(\nabla \hat{\varphi}_{c_1}^n, \nabla \hat{\varphi}_{c_2}^m). \quad (4.16)$$

The orientation similarity is given by dot product

$$\text{Ang}(v_1, v_2) = \frac{|v_1 \cdot v_2|}{\|v_1\| \|v_2\|}, \quad (4.17)$$

and the magnitude similarity by a Gaussian function with standard deviation of a third of the maximum frequency

$$\text{Mag}(v_1, v_2) = \exp\left(-\frac{1}{2} \frac{(\|v_1\| - \|v_2\|)^2}{(\pi/3)^2}\right). \quad (4.18)$$

As shown in Figure 4.1, when compared to the color-processing algorithm of Gastal and Oliveira (GASTAL; OLIVEIRA, 2017), our multi-channel optimization leads to fewer color distortions when used for spatial frequency adjustments. This occurs because their

Figure 4.1 – Our multi-channel phase unwrapping leads to a frequency-adjusted image, in (b) and (e), that is faithful to the colors of the original photograph, in (a) and (d). The remapping algorithm of Gastal and Oliveira (GASTAL; OLIVEIRA, 2017), in (c) and (f), reduces frequencies but is not able to correctly preserve colors. Adjustment of $\alpha = 0.2$ in both images.

Source: (a) Original photograph © Monceau on Flickr. Used with permission. (b, d, e) The authors. (c, f) (GASTAL; OLIVEIRA, 2017)

wave-alignment step is limited to processing a single color channel (either the lightness channel or one PCA-decorrelated channel). Our approach, on the other hand, allows for a joint optimization that preserves inter-channel correlations, resulting in improved frequency adjustment results. As such, we recommend the use of our multi-channel optimization in RGB space (Eqs. (4.12), (4.14) and (4.15)) when processing images which contain cross-channel frequency correlations.

For processing grayscale images one should use Eq. (4.12) directly. Similarly, for color images which contain most of their high-frequency contrast variations in the lightness channel, the combination of our phase unwrapping technique with the PCA-guided processing of Gastal and Oliveira (GASTAL; OLIVEIRA, 2017) leads to the best results. In practice, this means extracting and unwrapping the phase of only one “channel,” represented by the principal direction of color variation in RGB space. This procedure was used for generating the results shown in Figure 6.3, where the regions with detected high frequencies are mostly monochromatic (bluish gray and dark blue, respectively). We visually inspected the image to determine if the content is mostly monochromatic,

alternatively we may also compare the remapped high-frequency content.

5 IMPLEMENTATION DETAILS

Eq. (4.12) or (4.14) is solved once in a pre-processing step and the phase values $\varphi^n(p)$ are stored for subsequent real-time frequency adjustment. The resulting linear system is sparse and positive definite, which we solve using a QR factorization from SuiteSparse (DAVIS, 2011). We have also experimented with an iterative conjugate-gradient solver, but it was slower than the direct solver and stopping the iterations before convergence leads to visual artifacts during frequency adjustment (misalignment between structured patterns due to incomplete phase unwrapping). Frequency adjustment by a factor $\alpha > 0$ is obtained simply by replacing θ^n with $\alpha\varphi^n$ in Eq. (3.3). The resulting frequency-adjusted image is then reconstructed as

$$h_\alpha = s_\alpha + d, \quad \text{where} \quad s_\alpha(p) = \sum_{n=1}^{N(p)} B^n(p) \cos(\alpha\varphi^n(p)). \quad (5.1)$$

In practice, however, it is best to compute $s_\alpha(p)$ using a neighborhood Ω_p around p . This is achieved through the use of a weighted reconstruction window w , centered at p and covering Ω_p . This guarantees a smooth reconstruction for s_α , even in the presence of noisy measurements or missing phase values. The corresponding reconstruction equation, with windowing, becomes

$$s_\alpha(p) = \sum_{q \in \Omega_p} \sum_{n=1}^{N(q)} w(p-q) B^n(q) \cos(\alpha\{\varphi^n(q) + \nabla\hat{\varphi}^n(q) \cdot (p-q)\}). \quad (5.2)$$

The expression $\varphi^n(q) + \nabla\hat{\varphi}^n(q) \cdot (p-q)$ is a first order (linear) extrapolation of the phase $\varphi^n(q)$ from q to p , where $p \cdot q$ is the dot product between the points p and q . The appropriate reconstruction window $w(x, y)$ depends on the parameters of the Gabor decomposition being employed for space-frequency analysis, as discussed in detail by Gastal and Oliveira (GASTAL; OLIVEIRA, 2017). In Eq. (5.2), w is given by the square of an L^2 -normalized Gaussian, with the same standard deviation as the Gabor atoms.

5.1 Sparse Sampling

One can improve the computational and storage costs of the pre-processing step by performing a subsampled Gabor decomposition. For a properly selected subsampling

step, this procedure conveys the same information as the fully-sampled alternative. This is possible because a fully-sampled Gabor space is highly redundant (MALLAT, 2009). Thus, in practice, one can compute the gradient measurements $\nabla \hat{\varphi}(p)$ only every τ pixels (in all directions). This reduces the computational and storage costs by a factor of τ^2 when processing images, and τ^3 for videos. The maximum admissible value of the stepsize $\tau \geq 1$ depends on the Gabor atom's spread (MALLAT, 2009). Note that the windowed reconstruction from Eq. (5.2) naturally handles subsampling: any pixel q without phase information has $N(q) = 0$, and its computed final value $s_\alpha(q)$ is obtained from the neighborhood Ω_q . In this case, however, the proper reconstruction window w is not Gaussian and depends on τ . A formula for obtaining w can be found in (GASTAL; OLIVEIRA, 2017). One should also adapt the integration rule in Eq. (4.3) to consider the step size τ by dividing the left-hand-side by τ , and also computing the shift operators T_τ using shifts of τ (SOLOMON, 2015).

5.2 Use with Image Resizing

Frequency adjustment is orthogonal to image resizing. For reducing the number of pixels (downscaling), two procedures exist.

Procedure 1: the frequency-adjusted image ($h_\alpha = s_\alpha + d$ in Eq. (5.1)) is first reconstructed at the original resolution (preferably with windowing, Eq. (5.2)), and subsequently downsampled using a suitable strategy, which often involves a prefilter to avoid aliasing (NEHAB; HOPPE, 2014).

Procedure 2: the residual image (d in Eq. (5.1)) is downsampled independently (by a suitable strategy), and the frequency-adjusted detail (s_α) is reconstructed *directly at the target resolution*. That is, for a downscaling factor $r < 1$, the values of $s_\alpha(p)$ are sampled every $1/r$ units. Note that Eq. (5.2) naturally supports such fractional sampling in its argument p (with the understanding that the pixels q , inside the summations, represent only valid pixels from the original image's resolution, i.e., located at integer coordinates). The linear extrapolation in the cosine's argument computes the correct phase for any fractional location p , based on its shift $p - q$ from q . The extrapolation uses Euler's method to compute the value of the phase from the frequency. Note that as the wave is a planar wave, the second and higher derivatives of the phase are 0 and the error of the method is 0. For this strategy to be effective, however, the adjustment factor α must take into account the target resolution, as it defines a maximum representable spatial frequency; otherwise,

fractional sampling will be prone to aliasing (see Section 5.2.1, below).

For image upscaling ($r > 1$), one must reconstruct the high-frequency details s_α directly at the target (higher) resolution. Our implementation uses the GPU’s mip-map hardware for filtering and rescaling the residual image d , and uses a pixel shader to compute the frequency-adjusted detail s_α , at the target resolution using fractional sampling (either downsampled or upsampled). For additional flexibility, our implementation allows for independent control of frequency adjustment (α) and image resizing (r) factors, with instant feedback on the resulting image. As one would expect, depending on the selected values for these parameters, aliasing may occur. Alternatively, the two factors may be “linked,” automatically adjusting the frequencies in synchrony with image resizing and avoiding aliasing.

For general purpose, procedure 2 should be used if it is possible to guarantee that the adjusted frequencies are representable in the reconstructed space (see Section 5.2.1), otherwise the procedure able to represent the highest frequency in the reconstruction space should be used (i.e. if $r > 1$ use procedure 2 otherwise use procedure 1). If a special resizing algorithm is being used, such as an edge-aware filter or a deep-learning based strategy, then procedure 1 should preferred.

If the focus is performance, then the procedure with the least amount of samples in the reconstruction space should be used as long as the adjusted frequencies are representable in the reconstructed space. If the frequencies are not representable then a low-pass filter should be used after reconstructing the high-frequency content, furthermore, if downscaling procedure 1 is recommended as the downscaling algorithm should remove the generated aliasing.

Figure 5.1 illustrates the decomposition $s_\alpha + d$ (Eq. (5.1)) for the image in (a), subject to frequency adjustment using $\alpha = 0.25$. It also shows the unwrapped phase components (for the red channel) used by our application for real-time frequency adjustment. It is important to note that the residual d may be upsampled or downsampled using any applicable strategy before being added to the reconstructed frequencies in s_α . As such, our method is orthogonal to any resampling technique (Chapter 6).

5.2.1 Choosing α

The selection of the frequency-adjustment factor α should consider some criteria. Thus, let $\omega_{\max} \in [0, \pi]$ be the maximum frequency (in *radians per pixel*) present in an

image at its original resolution. After undergoing simultaneous resizing by a factor $r > 0$ and frequency adjustment by a factor $\alpha > 0$, $\omega_{\max}$ becomes $\omega_r = (\alpha/r) \omega_{\max}$. In order to avoid aliasing in the resulting image, the Nyquist limit must be observed:

$$(\alpha/r) \omega_{\max} \leq \pi. \quad (5.3)$$

If the original image encompasses the full available bandwidth, then by definition $\omega_{\max} = \pi$, and Eq. (5.3) results in the requirement that $\alpha \leq r$. Most commonly, however, structured patterns in images rarely represent the highest possible frequency (not considering the piecewise discontinuities that occur, for example, along object edges—the frequencies associated with these discontinuities are deliberately not included in the high-frequency detail s by the space-frequency analysis). As such, one can often select $\alpha > r$ and still obtain good (aliasing-free) results. The exact value of α used for a particular application is a subjective choice, left to the user, but should satisfy Eq. (5.3) to avoid aliasing.

In our implementation, we provide the following options: (i) α freely controllable by the user; (ii) $\alpha = r\alpha^*$ where α^* is user-controllable (i.e., α is proportional to r with some user-adjustable proportionality constant); and (iii) $\alpha = \sqrt{r}\alpha^*$. We call this last (empirically-defined) option a “perceptual” selection for α , as it leads to downsampled images that better match our expectations (refer to Figure 5.2). However, no user validation was performed to validate this claim.

5.3 Processing Videos and 3-D Volumes

Extending our technique to videos is straightforward. While one could consider independently applying our method to the individual frames, this introduces visual artifacts due to discontinuities in the recovered phases for adjacent frames. A more robust solution processes the whole 3-D (x, y, t) video volume at once. For the phase unwrapping pre-processing step, this is done by adding a new set of linear constraints relating to the third dimension: that is, simply by considering $\circ \in \{x, y, t\}$ in Eqs. (4.8) and (4.12) and all associated equations. The resulting 3-D phase functions $\varphi^n(x, y, t)$ are stored and later employed for real-time frequency adjustment by combining the frequency-adjusted components $\cos(\alpha \varphi^n(x, y, t))$ according to Eq. (5.2) (windowed reconstruction with a 3-D neighborhood Ω_p). As in the case for images, this process is completely independent per pixel and can be computed in parallel on a GPU. Furthermore, this can be done

during video playback where only the current frame has to be reconstructed. The same considerations apply for processing 3-D volumes in general (3-D scalar fields).

5.4 Splitted phase unwrapping for anisotropic frequency adjustment

To perform *anisotropic frequency adjustment*, that is frequency adjustment with different factors for each direction such as the one in Figure 5.3, one should project the phase gradient vectors in some orthogonal basis of the working space ($\mathbb{R}^2$ for images and $\mathbb{R}^3$ for videos) and integrate the projected phase gradients independently, obtaining a phase scalar field associated with each vector in the basis. For simplicity, it is desirable to work with the standard basis, as it is an orthogonal basis and because digital signals are sampled discretely along the standard axes.

Each recovered phase field encodes the frequency in the basis directions, thus, by using a scaling factor for each direction, one obtains anisotropic frequency adjustment. The phase can thus be seen as a vector ($\varphi = [\varphi_x, \varphi_y, \dots]$) and frequency adjustment as a functional applied to this phase vector.

The reconstruction process of phase vectors is very similar to phase scalars, the only difference is that a linear mapping is applied to obtain a scalar:

$$s_\alpha(p) = \sum_{n=1}^{N(p)} B^n(p) \cos(\alpha \cdot \varphi^n(p)), \quad (5.4)$$

where $\alpha = [\alpha_x, \alpha_y, \dots]$ is a vector of frequency-adjustment factors of the corresponding directions (which operates as a functional) and φ^n are the phase vectors.

The main motivation to perform anisotropic frequency adjustment is when a different downscaling factor is being used for each dimension. For instance, Figure 5.4 shows an example of anisotropic frequency adjustment with all power of two combinations of horizontal and vertical scaling factors. Another example of application is for video downscaling, as usually the spatial and temporal content are downscaled independently and potentially with different downscaling factors, e.g. the frames are usually subsampled by an integer factor while the spatial content is downscaled by some low-pass filter.

5.5 Complexity Analysis

The cost of the space-frequency analysis (Section 2.2.4) for an input image with $K = w \times h$ pixels using an $L \times L$ window is $O(KL \log L)$ due to the fast Fourier transform. Phase unwrapping requires solving a sparse linear system with $O(K/\tau)$ non-zero elements, where τ is the sampling step.

To reconstruct a target image with M pixels and at most N detected phase components per pixel, it requires MN element-wise multiplications of an $L \times L$ reconstruction window by the (per-pixel) cosine functions corresponding to the unwrapped phases, resulting in a cost of $O(MNL^2)$. Using sparse sampling with a step τ in both dimensions, the cost reduces to $O(MN \lfloor L/\tau \rfloor^2)$.

The reconstruction algorithm uses as input a residual image with the same dimensions $w \times h$ as the original image, and 2-D arrays containing pre-computed phase (unwrapped), amplitude, and gradients for the N components. Each such array has dimensions $w + L \times h + L$. The extra L rows and columns are required as reconstruction windows outside the image may affect pixels inside it. When sparse sampling with a step τ , only $\lfloor (w + L)/\tau \rfloor \times \lfloor (h + L)/\tau \rfloor$ samples are stored for each array, in which case the amount of memory required by the reconstruction algorithm is $O(wh + M + 4N \lfloor (w + L)/\tau \rfloor \lfloor (h + L)/\tau \rfloor)$.

Figure 5.1 – Decomposition used for frequency-adjustment of the reference image in (a). (b) Residual returned by the space-frequency decomposition, containing the low-frequency information of the input image. Note the absence of the stripes on the fabrics when compared to (a). (c) High-frequency detail s_α , after being frequency-adjusted with $\alpha = 0.25$ by our technique. Note that values in s_α are permitted to be negative, and the visualization in (c) maps all values to $[0, 1]$ for viewing purposes only. (d) Unwrapped phases (in radians) obtained by solving the multichannel Eq. (4.14), used to compute s_α in (c) (only components for the red channel are shown).

(a) Reference image (720×576)

(b) Residual image (d)

(c) Frequency-adjusted detail (s_α)

(d) Unwrapped phase components for the red channel ($\varphi_R^1, \varphi_R^2, \varphi_R^3$)

Source: (a) “Barbara” image attributed to Allen Gersho (public domain). (b–d) The authors.

Figure 5.2 – Perceptual frequency adjustment. The image on the left was downsampled using factors $r \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}\}$. If no frequency adjustment is performed before downsampling (**right-bottom row**), the high-frequency information is lost at the smaller image sizes. If frequencies are adjusted using a factor $\alpha = r$ (**right-top row**), the striped pattern is correctly preserved at all resolutions. However, note how all resulting images have sinusoids with the same period, measured in pixels. This results in a “perceptual dissonance” since, while the image size is being reduced, the pattern is perceived as staying the same size (since the wave periods are constant). This is solved by using a “perceptual” adjustment factor $\alpha = \sqrt{r}$ (**right-center row**), which reduces the pattern’s period lengths while still preserving the stripes.

Source: The authors.

Figure 5.3 – Example of anisotropic frequency adjustment. The grid shows how anisotropic frequency adjustment works on a chirp, a signal with increasing frequency, defined by $\cos(2\pi(x^2 + y^2))$, where the coordinate systems origin occurs in the center of the image. Each row and column has a horizontal (α_h) and vertical (α_v) adjustment factor, respectively. The diagonal highlighted with a *red border* illustrates the isotropic case, where $\alpha_h = \alpha_v$. The center of the chirp contains low frequency content that is not adjusted, and the anisotropic case shows how the radial chirp becomes an ellipsoid as the semi-axes are scaled by the α_h/α_v factors, with the extreme case $\alpha_h = 0/\alpha_v = 0$ ending in a planar vertical/horizontal chirp. Note that when $\alpha_h = \alpha_v = 0$ the wave becomes a constant function, as the frequency becomes 0. Please inspect the images using a proper zoom to avoid aliasing generated by the PDF Viewer.

Source: The authors.

Figure 5.4 – Example of application for anisotropic frequency adjustment. The image shows how anisotropic frequency adjustment works on the picture of “Barbara”. On the top-left we have the picture in its original resolution and to its immediate right and below it the image has its width and height downsampled to half the original size, respectively. The process is repeated downscaling the original image to half the resolution of the neighbors. For each downsampled image the horizontal and vertical frequencies are adjusted independently by the downsampled width and height divided by the original width and height, respectively. Notice how the waves orientation is preserved after performing non-uniform downscaling.

Source: Original “Barbara” image attributed to Allen Gersho (public domain). Downsampled versions generated by the authors.

6 RESULTS

We have applied our frequency-adjustment technique to a large number of images, including photographs, hand drawings, as well as synthetic ones. Here we demonstrate the potential and flexibility of our method through a few representative examples illustrating both increase and decrease of spatially-localized frequencies in combination with image resizing, resampling, and super-resolution techniques. Space-frequency factorization and phase unwrapping take a few seconds for a full HD image and are performed as a pre-processing. For instance, the space-frequency factorization of the $1,618 \times 2,048$ -pixel image shown in Figure 1.1 (a) takes about 8.8 seconds (detecting the relevant waves: 7.83 sec; building and solving an $83,633^2$ linear system: 0.93 sec and 0.04 sec, respectively). These measurements were made on a 6-core 3.2 GHz CPU. For all examples shown in the thesis, frequency adjustment and image reconstruction are performed in real time on a pixel shader. Our technique performs frequency adjustments on the 3 megapixel image shown in Figure 1.1 (a) at about 70 fps using $\tau = 10$ on an AMD R9 280X GPU. Although spectral remapping can also detect relevant phases as some pre-processing, it requires building and solving a linear system for phase alignment every time the target resolution changes. For the image shown in Figure 1.1, these operations take about 1.5 seconds on a 6-core 3.2 GHz CPU.

In all comparisons with other techniques, we used software provided by their own authors. Our technique can generate its results directly at the target image resolution. Nevertheless, in all downscaling examples shown in the thesis we perform frequency adjustment at the same resolution as the input image, followed by a resampling to the target dimensions. We do so to stress the fact that our method is orthogonal to and complements existing downscaling techniques (KOPF; SHAMIR; PEERS, 2013; NEHAB; HOPPE, 2014; ÖZTIRELI; GROSS, 2015; SACHT; NEHAB, 2015; WEBER et al., 2016), we also stress that our method should be used to preserve the high-frequency content’s direction at the cost of changing the high frequencies magnitude. Thus, along the following examples we will be using different resampling techniques. We encourage the reader to inspect our results by zooming in on the full-resolution images embedded on the PDF.

Figure 1.1 (a) shows a hand drawing containing regions with very high spatial frequencies. Figure 1.1 (e) shows the downscaling of the reference image to 135×171 pixels using our technique, taking 0.01 seconds for frequency adjustment and image reconstruction. Note how the structured details of the high-resolution reference image in (a) are

Figure 6.1 – Examples of adjusting the frequency of complex natural patterns for downscaling. (a) Reference high-resolution images. (b) Downscaled versions of the images in (a) generated by performing frequency adjustment with our technique, followed by resampling using a Lanczos kernel. (c) Downscaled versions of the images in (a) using Lanczos resampling only. Note how our result in (b) manages to preserve the structure in the sunflower’s disk and the honeycomb’s tiling, which are otherwise discarded by standard downscaling prefilterers (c).

(a) Reference (left: 1920×1278 , right: 1920×1440)

(b) Ours + Lanczos

(c) Lanczos

(288×192 , $r = 0.15$, $\alpha = 0.45$)

(288×192)

(288×216 , $r = 0.15$, $\alpha = 0.6$)

(288×216)

Source: Sunflower image courtesy of Uschi Dugulin and Honeycomb image courtesy of Vijaya Narasimha. (b–c) The authors.

Figure 6.2 – Stencil masks used to apply independent scaling factors to the spatial frequencies of different image regions: pants (red), scarf (blue), and table cloth (green). These stencils were created using GIMP’s *Intelligent Scissors* tool to segment the corresponding regions, which were further dilated by a few pixels to guarantee that no pixels at the borders of high-frequency regions were left out. Alternatively, such regions could potentially be detected by our technique as clusters of neighboring pixels whose reconstructing cosine functions have similar amplitude and phase values.

Source: The authors.

preserved by our technique at the lower image resolution in (e), which are completely removed by the Cubic O-MOMS filter in (g). By comparison, spectral remapping (GASTAL; OLIVEIRA, 2017), in (f), generates a similar result but is two orders of magnitude slower to compute. In both cases, the resampling to the target resolution was performed using the Cubic O-MOMS filter (NEHAB; HOPPE, 2014). Figures 1.1 (b–d) show examples of frequency-adjusted results using our technique for different adjustment factors α . We note that the background circle uses a form of dithering (black dots on a white background), and perhaps the artist’s objective was for it to look “grayish” when observed from a certain distance. If that was the original intent, then a standard linear blurring filter should be preferred over our frequency-adjustment operation, as the high-frequency content should not be preserved in this case.

Figure 6.1 (a) shows two photographs containing natural high frequency patterns. By reducing the frequency of these structures using our approach, one is able to preserve their appearance in a much lower resolution, as shown in (b) for a scaling factor of $0.15\times$. For these examples the frequency-adjusted images were reconstructed at the original resolution and subsequently downsampled using a 7-lobed Lanczos kernel. If one downscales the reference image using this same kernel, but without adjusting frequencies using our approach, the detailed patterns of the sunflower and honeycomb are lost, as

shown in Figure 6.1 (c).

Although frequency adjustment is performed at pixel level, it potentially affects content over the entire image. Thus, additional creative control is provided by allowing the user to specify a different frequency scaling factor to individual portions of an image, or to delimit the set of pixels the frequency adjustment should be applied to. This can be achieved by providing a stencil mask specifying the target pixels (Figure 6.2). The mask does not need to precisely delimited the region between the high and low frequencies, however it may reduce color bleeding from the reconstruction windows (see Section 7.1). In regions where there are two high frequency waves that should be adjusted with different factors a more precise delimitation is required, specially for big resizing factors.

Figure 6.3 shows the results of our technique applied to some textured fabrics. In these examples, we demonstrate the combined use of our method with Lim et al.’s. (LIM et al., 2017) super-resolution technique. This illustrates our method’s ability to complement super-resolution strategies to obtain high-resolution images with a wider frequency bandwidth. Figures 6.3 (b) and (h) show downsampled versions of the corresponding reference images to 80×60 and 75×75 pixels, respectively, produced by our method. In both cases, resampling was performed using the technique of Weber et al. (WEBER et al., 2016). Figures 6.3 (c)/(i) show results generated by our method when increasing the frequencies by a factor $\alpha = 2$, while preserving the original resolution ($r = 1$). Note how the spatial frequencies in (c)/(i) are higher than the ones in the corresponding reference images (a)/(g). They are also higher than the ones in the results produced by the super-resolution technique of Lim et al. (LIM et al., 2017), shown in (e)/(k). Figures 6.3 (d)/(j) show the results obtained by combining our method with Lim et al.’s. This hybrid solution consists of applying the super-resolution technique to the residual image, which is then added to the adjusted high-frequency detail, according to Eq. (5.1). Note that such a procedure can be used with any super-resolution method. Figures 6.3 (d)/(j) show fabrics with significantly higher frequency content in comparison to the plain super-resolution results in (e)/(k), and also to standard bicubic upsampling in (f)/(l). These results demonstrate the effectiveness and simplicity of such a hybrid approach.

Figure 6.4 shows a photograph of a lizard exhibiting high-frequency details. Figure 6.4 (b) shows the downsampled version of the reference image to 220×293 pixels generated with frequency adjustments by our technique. Resampling to the target resolution was performed with the optimized quasi interpolator OQI3 (SACHT; NEHAB, 2015). Figures 6.4 (c) and (d) show results produced by our method after scaling the frequencies

Figure 6.3 – Examples of downscaling and upscaling with frequency adjustments. (a) and (g) Reference images. (b) and (h) Downscaled versions of (a) and (g) generate by performing frequency adjustment with our technique, followed by resampling using the method of (WEBER et al., 2016). (c) to (d) and (i) to (j) show examples of frequency adjustments using our technique. (c) and (i): Increasing spatial frequency using a scaling factor $\alpha = 2$, while keeping the dimensions of the reference image. (d) and (j) Increasing the spatial frequency using $\alpha = 2$ while resizing the image by $2\times$ (using Lim et al.’s (LIM et al., 2017) super-resolution method for upscaling the residual). (e) and (k) Upscaling by $2\times$ using only Lim et al.’s (LIM et al., 2017) super-resolution method. (f) and (l) Upscaling by $2\times$ using standard bicubic upsampling. Note how our approach is able to generate high-frequency detail that is significantly different than existing classic and state-of-the-art methods.

Source: (a, g) Original Images from MSRA Salient Object Set. (b–f, h–l) The authors.

by factors of $\alpha = 0.25$ and $\alpha = 1.25$, respectively. One should compare the details on the belly and legs of the lizards with the ones on the reference image. For the examples in Figures 6.4 and 6.5 no comparisons with the super-resolution technique of Lim et al. (LIM et al., 2017) are provided (the available implementation did not handle the sizes of the reference images).

The striped shirt in Figure 6.5 (a) contains a high-frequency pattern, being a challenging test case for image downscaling techniques. Figure 6.5 (b) compares the results of our method (top) with spectral remapping (bottom) for the case of an aggressive downscaling of the reference image to 200×150 pixels. In both cases, the actual downscaling step was performed with Lanczos filtering followed by resampling using a cubic B-spline. Figure 6.5 (c) shows an example of frequency increase (applied to an already high-frequency pattern). For this result, we kept the dimensions of the original image ($r = 1$) while increasing the frequency of the stripes using $\alpha = 1.5$. This example also demonstrates the effectiveness of our solution for processing color images (Section 4.3), as misalignment in the phases of the individual color channels would introduce noticeable artifacts in such a large image (Figure 4.1 (f)).

Figure 6.6 illustrates the use of our technique for processing a video sequence. As discussed in Section 5.3, the whole 3-D (x, y, t) volume is processed simultaneously, thus preserving temporal coherence of the frequency-adjusted details.

Figure 6.4 – Picture of a lizard (courtesy of William Warby). (a) Reference image. (b) Downscaled version of (a) generated with (top) and without (bottom) our technique for frequency adjustment and both resampled to 220×293 using the method of Sacht and Nehab (SACHT; NEHAB, 2015). (c) and (d) Examples of spatial frequency adjustment generated by our technique. The images were kept at the original resolution ($r = 1$): (c) Frequency decreasing obtained with a scaling factor $\alpha = 0.25$. (d) Frequency increasing obtained with a scaling factor $\alpha = 1.25$. Note the adjusted frequencies of the details on legs and belly of the lizard.

Source: (a) Picture of a lizard courtesy of William Warby. (b–d) The authors.

Figure 6.5 – Shirt. (a) Reference image. (b) Downscaled versions of (a): (top) generated with our technique and (bottom) obtained with spectral remapping (GASTAL; OLIVEIRA, 2017). In both cases, resampling to 200×150 pixels was done applying Lanczos filtering followed by resampling using a cubic B-spline. (c) Result produced by our techniques with same resolution as the reference image ($r = 1$), but with spatial frequency increased by a factor $\alpha = 1.5$. Please zoom-in to avoid aliasing artifacts caused by the PDF reader’s resampling filter.

Source: The authors.

Figure 6.6 – Example of video frequency adjustment. (a) The reference video contains high-frequency patterns on the jacket, which move around from frame to frame following the person’s movement. (b) Our frequency adjustment technique is applied to the whole 3-D space-time video volume (Section 5.3), reducing the frequency of the patterns and preserving them at the target (lower) resolution.

(a) Frame from the reference video (760×720)

(b) Frames from our frequency-adjusted video (downscaled to 228×216 using Lanczos filtering)

Source: (a) Original video (public domain) from <https://media.xiph.org/video/derf/>.
(b) The authors.

7 THEORETICAL DISCUSSION

In this chapter we present discussions on relevant theoretical details.

7.1 Limitations

As discussed in Section 2.2.4, we use a Gabor space-frequency decomposition for measuring the instantaneous frequencies $\nabla \hat{\varphi}$ present in the image’s spectrum (GASTAL; OLIVEIRA, 2017). One limitation of this approach is that the (Gaussian-windowed) Gabor atoms have a fixed bandwidth in space and frequency, controlled by a user-specified frequency threshold $\bar{\omega}$. If one wants to detect the instantaneous frequencies of very slowly varying components, the threshold $\bar{\omega}$ must be small in order to obtain a good frequency (spectral) resolution. But in turn this comes at the cost of poorer spatial resolution, since space and frequency locality are conflicting requirements (MALLAT, 2009). In practice, this means that reducing the frequency threshold leads to larger analysis and reconstruction windows for the Gabor frame. Such large windows may cause some detected high-frequencies to bleed into nearby pixels during frequency adjustment. These “blended frequencies” are often hidden by the downscaling procedure, but may be visible when the resulting image is upsampled or has its size preserved. This can be seen, for example, around the letters in Figure 6.3 (d) and around the buttons in Figure 6.5 (c). This problem could potentially be mitigated by replacing the Gabor step with a Wavelet strategy for detecting instantaneous frequencies (DELPRAT et al., 1992). Since wavelets adapt their spatial support based on the frequency bands of the signal, they define a multiscale approach which would use the smallest possible window for each detected frequency.

As mentioned in Section 4.3, for some color images it is best to use the PCA-based approach for phase unwrapping (processing only a single “channel”). This occurs because, in some situations, the inter-channel alignment is not enforceable without breaking intra-channel phase constraints. This is the case of the images in Figure 6.3 (a) and (g), which are highly monotone, and for which a multi-channel RGB alignment introduces some small chromatic distortions (Figure 7.1).

Figure 7.1 – Chromatic distortions might happen when enforcing multi-channel alignment for images that are mostly monotone (see the button on the left and the shirt texture on the right). Images obtained using $\alpha = 0.5$ and $r = 1$.

Source: The authors.

7.2 Comparison to Spectral Remapping

In this section we will discuss the theoretical differences and relations between our frequency adjustment and the spectral remapping method (GASTAL; OLIVEIRA, 2017).

Spectral remapping uses the fact that waves are locally planar to reconstruct an image as $A \cos(ax + by + c)$. Note that, as the wave is planar, its frequency is constant, thus the phase varies linearly with the x and y coordinates according to the horizontal and vertical frequencies a and b . The phase shift is computed as the constant that minimizes the pointwise difference between the windowed wave with respect to the phase-shift and the original image also windowed. This is equivalent to solving an initial value problem to find the integration constant from integrating the detected planar wave frequencies.

This means that spectral remapping also computes a phase like frequency adjustment, however the phase-shift factor is computed modulus 2π , therefore after remapping the wave the recovered phase has the same discontinuities problem as applying frequency adjustment to a wrapped phase field. The spectral remapping solution for this is to use the fact that the remapped frequency is correct so that only a set of phase-shift factors that *aligns* the waves needs to be found. This is done by minimizing the difference between the cosines of the remapped waves in overlapping regions, similar to the initial value problem used to find the original phase-shift, but now with remapped frequencies.

This solution is what distinguishes the two methods, while spectral remapping works with waves with remappable frequencies that are shifted by a certain set of phase-

shifts that aligns the color of overlapping waves. Frequency adjustment works with a decomposition in another space, the *phase space*, which relates to color space by an amplitude modulated cosine function and is smooth so that the recovered signal in color space is continuous even after applying linear transformations in phase space.

In summary, the color of the remapped waves dictates the alignment, i.e. continuity of colors, in spectral remapping, while, in frequency adjustment, the smoothness of the phases guarantees that the waves are aligned in color space.

7.3 Future Works

Given the novelty of this area several improvements and research can be done, thus in this section we highlight the most significant contributions that could be introduced in future works.

It is possible to apply frequency adjustment to images with aliasing to obtain a version with frequency adjustment. For this, note that aliasing is caused by high frequencies that are not representable in the lower resolution, but still present in the spectrum. The non-representable frequencies are *folded* into the representable range, thus if it is possible to detect the *unfolded frequency* it is possible to integrate with our method. Unfortunately, space-frequency decomposition only allows us to detect the folded frequency. However, if a method capable of detecting unfolded frequencies is used, it would be possible to adjust frequencies of images with aliasing.

Separating the detected waves into several components that could be unwrapped independently may also be possible using a seeded method to group together all waves with high alignment factor μ . However, because the decomposition is local, it is possible that two waves change direction in different regions of an image, in particular if in a given location two waves with similar enough frequencies are in the same direction only a single wave with amplitude equals to the sum of the other two waves would be detected. For phase unwrapping this would not be a problem, because the frequency is the same so when integrating the recovered phase would still be correct. However, if the objective is to separate the phases into continuous components this is a problem, because the wave would need to be divided in two with amplitudes equal to the neighboring waves.

There are several motivations to separate the waves into several groups, for instance one could have more control over frequency content such that only a single frequency component is adjusted. Another use would be to upsample τ -sparsed phase and amplitude

fields to obtain a dense phase field such that each pixel would have a set of phase and amplitude components. For this the phase and frequency should be used to estimate the phase at intermediate pixels and the amplitude could use more usual interpolation methods, such as bilinear interpolation. The amplitudes should be pre-multiplied by the Gaussian window value of the interpolated pixel. The advantage of upsampling the phase and amplitude would be to perform per pixel reconstruction removing the reconstruction window from the reconstruction method. This would remove all windows size L from the reconstruction complexity in Section 5.5, however it would increase the use of memory and processing by removing the τ sparsing factors.

Multichannel phase unwrapping could also be improved by adding a phase difference term. The idea is to enforce the constant term from the integral to have the same difference between channels such that the combined waves would generate the same color after unwrapping, possibly removing the chromatic distortions present in some images.

Other than these improvements to phase unwrapping and frequency adjustment, several applications could also be explored. For example, transformations in phase space could be studied to identify the effect in color space. An interesting transformation that could be studied is image morphing, such that the phase, amplitude and residual are morphed separately with different interpolation functions with speed proportional to the frequency or different morphing algorithms altogether. The reason is that the human eye tends to be more sensitive to high frequency content and so a different algorithm for high frequency content may have better results.

8 CONCLUSION

We presented a real-time technique for interactive adjustment of the spatial frequencies in images. It supports both decreasing and increasing of frequencies with localized spatial control, and is orthogonal to image resizing and resampling strategies. Conceptually, our method consists of factoring the input image, at a pixel level, into a set of high-frequency waves (cosines) plus some residual. Frequency adjustment is achieved by scaling the phase values of the detected waves at each pixel by some factor α . Image reconstruction is then obtained (also at pixel level) by accumulating the cosines of the scaled phase values and adding the corresponding residual.

This thesis introduces two additional contributions: *(i)* a phase-unwrapping technique to recover the phases of several real waves, and *(ii)* a technique for processing color images for use with frequency adjustments. Our phase-unwrapping solution naturally handles the sign ambiguity associated with the unwrapping of real waves. To the best of our knowledge, ours is the first algorithm to solve this problem in 2-D. Our method for processing color images avoids color artifacts associated with the use of the PCA-based color reconstruction described in (GASTAL; OLIVEIRA, 2017). We demonstrated the effectiveness of our technique on a large number of images subject to downscaling, upscaling, as well with their dimensions preserved. We have also shown how to extend our method to videos and how to combine it with conventional super-resolution techniques to obtain novel super-resolution results. Furthermore, the contributions of this thesis were presented at the Eurographics 2021 conference and published as an article in the Computer Graphics Forum journal (Eurographics full papers proceedings) (GERMANO; OLIVEIRA; GASTAL, 2021).

By performing frequency adjustment with real-time feedback, our technique expands the range of artistic and technical possibilities for image and video processing. This should stimulate the development of new applications involving real-time exploration of visual content on small as well as on very large displays. This could include, for instance, real-time image downscaling to adapt content to thumbnail sizes or to smartphone displays. Similarly, one could combine our frequency-enhancement solution with super-resolution methods to produce upscaled images for exploration on ultra high-definition displays.

REFERENCES

- BARBHUIYA, A. J. I.; HEMACHANDRAN, K. Wavelet transformations & its major applications in digital image processing. **Intl' Journal of Engineering Research & Technology (IJERT)**, p. 2278–0181, 2013.
- CHILDERS, D.; SKINNER, D.; KEMERAIT, R. The cepstrum: A guide to processing. **Proc. of the IEEE**, v. 65, n. 10, p. 1428–1443, 1977. ISSN 1558-2256.
- DAVIS, T. A. Algorithm 915, SuiteSparseQR: Multifrontal Multithreaded Rank-revealing Sparse QR Factorization. v. 38, n. 1, p. 8:1–8:22, 2011. ISSN 0098-3500.
- DELPRAT, N. et al. Asymptotic wavelet and Gabor analysis: extraction of instantaneous frequencies. **IEEE Transactions on Information Theory**, v. 38, n. 2, p. 644–664, mar. 1992. ISSN 1557-9654.
- FREEDMAN, G.; FATTAL, R. Image and video upscaling from local self-examples. **ACM Trans. Graph.**, ACM, New York, NY, USA, v. 28, n. 3, p. 1–10, 2010. ISSN 0730-0301.
- GASTAL, E. S. L.; OLIVEIRA, M. M. Spectral remapping for image downscaling. **ACM Trans. Graph.**, v. 36, n. 4, p. 145:1–145:16, 2017.
- GERMANO, R. L.; OLIVEIRA, M. M.; GASTAL, E. S. L. Real-time frequency adjustment of images and videos. **Computer Graphics Forum**, v. 40, n. 2, 2021.
- GHIGLIA, D. C.; PRITT, M. D. **Two-dimensional phase unwrapping: theory, algorithms, and software**. [S.l.]: Wiley, 1998. ISBN 978-0-471-24935-1.
- GOLDSTEIN, R. M.; ZEBKER, H. A.; WERNER, C. L. Satellite radar interferometry: Two-dimensional phase unwrapping. **Radio Science**, v. 23, n. 4, p. 713–720, jul. 1988. ISSN 1944-799X.
- GUILLEMAIN, P.; KRONLAND-MARTINET, R. Characterization of acoustic signals through continuous linear time-frequency representations. **Proceedings of the IEEE**, v. 84, n. 4, p. 561–585, abr. 1996. ISSN 1558-2256.
- HERSZTERG, I.; POGGI, M.; VIDAL, T. Two-Dimensional Phase Unwrapping via Balanced Spanning Forests. **INFORMS Journal on Computing**, v. 31, n. 3, p. 527–543, jul. 2019. ISSN 1091-9856, 1526-5528. ArXiv: 1812.08277.
- HUANG, N. E. **Hilbert-Huang transform and its applications**. [S.l.]: World Scientific, 2014.
- HUANG, N. E. et al. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. **Proceedings of the Royal Society of London. Series A.**, The Royal Society, v. 454, n. 1971, p. 903–995, 1998.
- ITOH, K. Analysis of the phase unwrapping algorithm. **Applied Optics**, v. 21, n. 14, p. 2470–2470, jul. 1982. ISSN 2155-3165.
- KIM, H. et al. Task-aware image downscaling. In: **Proceedings of the European Conference on Computer Vision (ECCV)**. [S.l.: s.n.], 2018. p. 399–414.

KOPF, J.; SHAMIR, A.; PEERS, P. Content-adaptive image downscaling. **ACM Trans. Graph.**, ACM, New York, NY, USA, v. 32, n. 6, p. 173:1–173:8, nov. 2013. ISSN 0730-0301.

LI, Y. et al. Learning a convolutional neural network for image compact-resolution. **IEEE Transactions on Image Processing**, IEEE, v. 28, n. 3, p. 1092–1107, 2018.

LILLY, J. M.; OLHEDE, S. C. On the Analytic Wavelet Transform. **IEEE Trans. on Information Theory**, v. 56, n. 8, p. 4135–4156, aug. 2010. ISSN 0018-9448, 1557-9654.

LIM, B. et al. Enhanced deep residual networks for single image super-resolution. In: **IEEE CVPR Workshops**. [S.l.: s.n.], 2017.

MALLAT, S. G. **A Wavelet Tour of Signal Processing: The Sparse Way**. 3rd. ed. [S.l.]: Elsevier, 2009. ISBN 978-0-12-374370-1.

MERTZ, L. N. Speckle imaging, photon by photon. **Applied Optics**, v. 18, n. 5, p. 611–614, mar. 1979. ISSN 2155-3165.

MITCHELL, D. P.; NETRAVALI, A. N. Reconstruction filters in computer-graphics. **SIGGRAPH Comput. Graph.**, ACM, New York, NY, USA, v. 22, n. 4, p. 221–228, jun. 1988. ISSN 0097-8930.

NEHAB, D.; HOPPE, H. A fresh look at generalized sampling. **Foundations and Trends in C.G. and Vision**, v. 8, n. 1, p. 1–84, 2014.

OPPENHEIM, A. V.; SCHAFER, R. W. **Discrete-time signal processing**. [S.l.]: Pearson Higher Education, 2010.

ÖZTIRELI, A. C.; GROSS, M. Perceptually based downscaling of images. **ACM Trans. Graph.**, ACM, New York, NY, USA, v. 34, n. 4, p. 77:1–77:10, jul. 2015. ISSN 0730-0301.

PRASAD, L.; IYENGAR, S. **Wavelet Analysis with Applications to Image Processing**. Taylor & Francis, 1997. ISBN 9780849331695. Available from Internet: <<https://books.google.com.br/books?id=IJt5zQw3v14C>>.

REINHARD, E. et al. Color transfer between images. **IEEE Computer Graphics and Applications**, v. 21, n. 5, p. 34–41, jul. 2001. ISSN 1558-1756.

SACHT, L.; NEHAB, D. Optimized Quasi-Interpolators for Image Reconstruction. **IEEE TIP**, v. 24, n. 12, p. 5249–5259, 2015. ISSN 1057-7149.

SCIVIER, M. S.; FIDDY, M. A. Phase ambiguities and the zeros of multidimensional band-limited functions. **JOSA A**, v. 2, n. 5, p. 693–697, may 1985. ISSN 1520-8532.

SOLOMON, J. **Numerical algorithms: methods for computer vision, machine learning, and graphics**. [S.l.]: CRC Press, Taylor & Francis, 2015. ISBN 978-1-4822-5188-3.

SONG, S. M.-H. et al. Phase unwrapping of MR phase images using Poisson equation. **IEEE TIP**, v. 4, n. 5, p. 667–676, 1995. ISSN 10577149.

SUN, W.; CHEN, Z. Learned image downscaling for upscaling using content adaptive resampler. **IEEE Transactions on Image Processing**, IEEE, v. 29, p. 4027–4040, 2020.

TAKAJO, H.; TAKAHASHI, T. Least-squares phase estimation from the phase difference. **JOSA A**, v. 5, n. 3, p. 416–425, mar. 1988. ISSN 1520-8532.

WEBER, N. et al. Rapid, detail-preserving image downscaling. **ACM Trans. Graph.**, ACM, New York, NY, USA, v. 35, n. 6, p. 205:1–205:6, 2016. ISSN 0730-0301.

Yang, J. et al. Image super-resolution via sparse representation. **IEEE TIP**, v. 19, n. 11, p. 2861–2873, Nov 2010. ISSN 1941-0042.

ZEBKER, H. A.; LU, Y. Phase unwrapping algorithms for radar interferometry: residue-cut, least-squares, and synthesis algorithms. **JOSA A**, v. 15, n. 3, p. 586–598, mar. 1998. ISSN 1520-8532.