Math 994-002: Applied and Computational Harmonic Analysis, MSU, Spring 2020

Lecture 27 & 28: 2D Directional Wavelet Frames

April 21 & 23, 2020

Lecturer: Matthew Hirn

6.2.1 Dyadic Maxima Representation

Section 6.2.2 of A Wavelet Tour of Signal Processing

Let us now return to the analysis of pointwise singularities of signals f via the decay of $Wf(u,s)$ as $s \to 0$. Let ψ be a real valued wavelet, and recall that a wavelet modulus maxima is defined as a point (u_0, s_0) such that $|Wf(u,s_0)|$ is locally maximum at $u = u_0$.

All of the results regarding wavelet coefficient decay and the pointwise regularity of $f(t)$ (including, in particular, Theorems 5.5, 5.7, and 5.8) hold for dyadic wavelet semi-discrete frames by restricting $s = 2^j$ for $j \in \mathbb{Z}$. Let (u_0, j) be a modulus maxima point of $Wf(u,j)$, meaning that

$$\frac{\partial Wf}{\partial u}(u_0, j) = 0 \tag{81}$$

Since $Wf(u,j) = f * \bar{\psi}_j(u)$, $\bar{\psi}_j(t) = 2^{-j}\psi(-2^{-j}t)$ and

$$\frac{d}{dt}\bar{\psi}_j(t) = -2^{-j}2^{-j}\psi'(-2^{-j}t) = -2^{-j}\overline{\psi'}_j(t)$$

equation (81) is equivalent to

$$f * \overline{\psi'}_j(u_0) = 0$$

Figure 34 shows the dyadic wavelet transform of a signal and the corresponding wavelet modulus maxima.

Let Λ denote the wavelet modulus maxima of f:

$$\Lambda = \{(u,j) \in \mathbb{R} \times \mathbb{Z} : f * \overline{\psi'}_j(u) = 0\}$$

Recall that the dictionary $\mathcal{D}$ of a dyadic wavelet transform is:

$$\mathcal{D} = \{\psi_{u,j}\}_{(u,j)\in\mathbb{R}\times\mathbb{Z}}$$

The set Λ defines a sub-dictionary of $\mathcal{D}$:

$$\mathcal{D}_\Lambda = \{\psi_{u,j}\}_{(u,j)\in\Lambda}$$

Furthermore, the completion of the span of $\mathcal{D}_\Lambda$ defines a closed subspace $\mathbf{V}_\Lambda$ of $\mathbf{L}^2(\mathbb{R})$:

$$\mathbf{V}_\Lambda = \overline{\operatorname{span} \mathcal{D}_\Lambda}$$

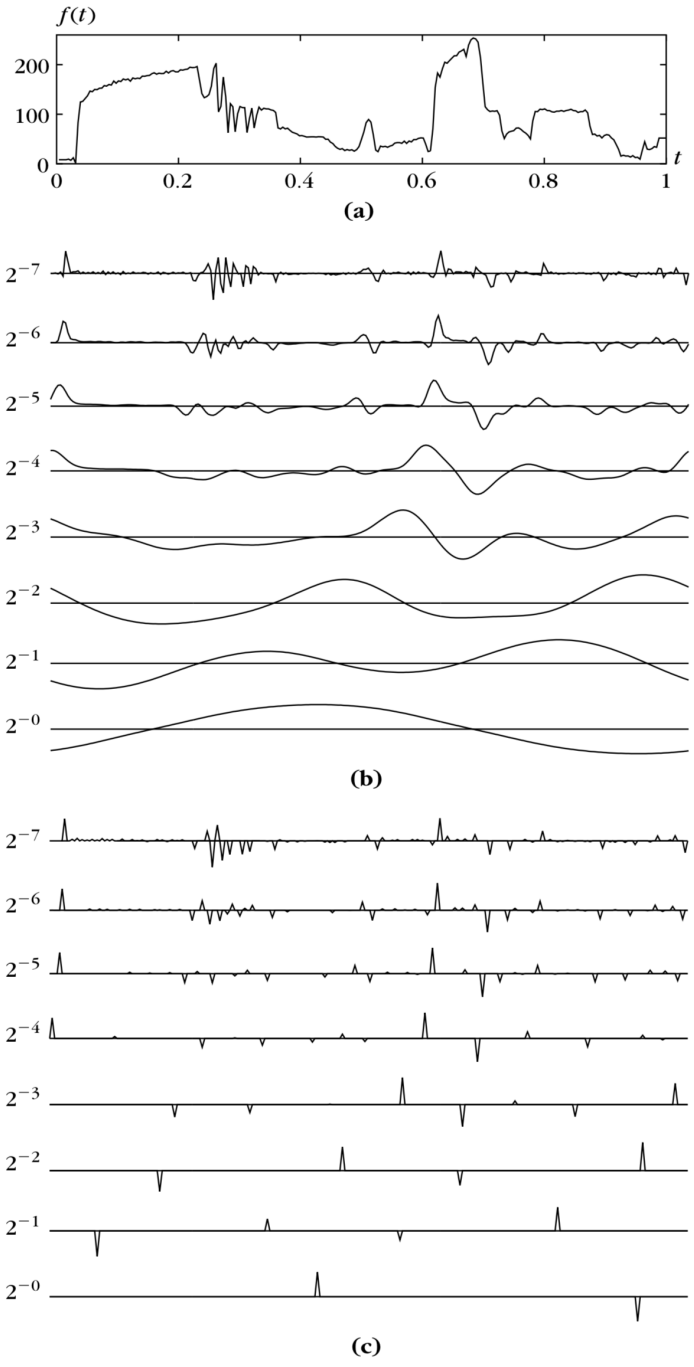

Figure 34: (a) The signal $f(t)$. (b) Dyadic wavelet transform computed with a wavelet $\psi = -\theta'$. (c) Modulus maxima of the dyadic wavelet transform.

We can therefore project f onto $\mathbf{V}_\Lambda$. Doing so amounts to computing an approximation f_Λ of f which is the signal synthesized from only the wavelet modulus maxima of f. It is computed with a dual synthesis as:

$$f_\Lambda = P_{\mathbf{V}_\Lambda} f = \sum_{(u,j)\in\Lambda} \langle f, \psi_{u,j}\rangle \widetilde{\psi}_{u,j}$$

For general dyadic wavelets, $f_\Lambda \neq f$. However, signals with the same modulus maxima differ from each other by small amplitude errors introducing no oscillations, so in numerical experiments $f_\Lambda \approx f$. If f is band-limited (meaning it has a compactly supported Fourier transform) and ψ is as well, then the wavelet modulus maxima define a complete representation of f and in this case $f_\Lambda = f$.

Figure 35 computes the projection f_Λ for the signal first introduced in Figure 34. The signal is not bandlimited, so the reconstruction is not perfect. However, Figure 35(b) shows that the approximation is quite good, and the relative error is approximately 2.5%. Figure 35 reconstructs the signal using only the top 50% of the wavelet modulus maxima coefficients. The sharpest signal transitions have been preserved, since they have the largest amplitude responses, however small texture variations are removed since the wavelet modulus maxima there have relatively small amplitudes. The resulting signal appears to be piecewise regular. In either case, we have achieve a lossy compression of f in which there is a large compression and the loss is not too large.

Exercise 73. Read Section 6.2.2 of *A Wavelet Tour of Signal Processing.*

6.3 Multiscale Directional Frames for Images

Section 5.5 of A Wavelet Tour of Signal Processing.

6.3.1 Directional Wavelet Frames

Section 5.5.1 of A Wavelet Tour of Signal Processing.

We now consider two dimensional wavelet semi-discrete frames for image analysis. Such semi-discrete frames are constructed with wavelets that have directional sensitivity, providing information on the direction of sharp transitions such as edges and textures.

Let $x = (x_1, x_2) \in \mathbb{R}^2$. A directional wavelet $\psi_\alpha(x)$ of angle $\alpha \in [0, 2\pi)$ is a wavelet having p directional vanishing moments along any one dimensional line of direction $\alpha + \pi/2$ in the plane but does not have directional vanishing moments along the direction α. The former condition means that:

$$\int_{\mathbb{R}} \psi_\alpha(\rho\cos\alpha - u\sin\alpha, \rho\sin\alpha + u\cos\alpha) u^k \, du = 0, \quad \forall\, \rho \in \mathbb{R},\ 0 \leq k < p$$

Such a wavelet oscillates in the direction $\alpha + \pi/2$ but not in the direction α.

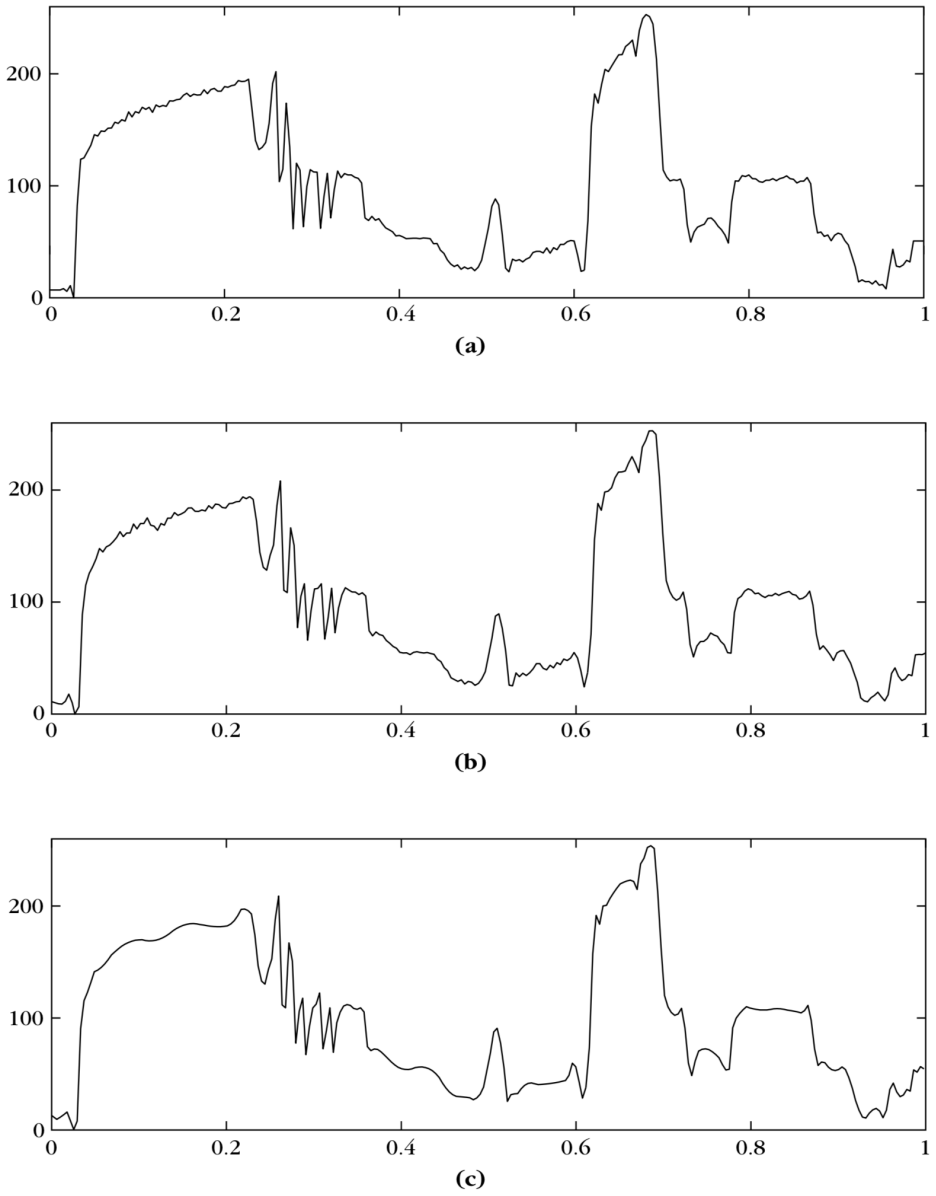

Figure 35: (a) The signal $f(t)$. (b) Signal approximation $f_\Lambda(t)$ using the dyadic wavelet modulus maxima shown in Figure 34. (c) Approximation recovered using only the largest 50% of the wavelet modulus maxima.

Let $\Theta \subset [0, \pi)$ denote the set of angles α. Typically Θ is a uniform sampling:

$$\Theta = \{\alpha = 2\pi k/K : 0 \leq k < K\}$$

The generators of a translation invariant dictionary are the dyadic dilations of each directional wavelet:

$$\{\psi_{j,\alpha}\}_{j\in\mathbb{Z},\,\alpha\in\Theta}, \quad \psi_{j,\alpha}(x) = 2^{-2j}\psi_\alpha(2^{-j}x)$$

Often the directional wavelets ψ_α are obtained by rotating a single mother wavelet ψ; we will come back to this shortly when we define two dimensional Gabor and Morlet wavelets. For real valued directional wavelets, Theorem 6.7 proves that the generating wavelets generate a semi-discrete frame if and only if there exists $0 < A \leq B < \infty$ such that

$$A \leq \sum_{j\in\mathbb{Z}} \sum_{\alpha\in\Theta} |\widehat{\psi}_\alpha(2^j\omega)|^2 \leq B, \quad \forall\, \omega \in \mathbb{R}^2 \setminus \{(0,0)\}$$

If the generating wavelets ψ_α are complex valued analytic wavelets, then they generate a semi-discrete frame if and only if

$$2A \leq \sum_{j\in\mathbb{Z}} \sum_{\alpha\in\Theta} |\widehat{\psi}_\alpha(2^j\omega)|^2 + \sum_{j\in\mathbb{Z}} \sum_{\alpha\in\Theta} |\widehat{\psi}_\alpha(-2^j\omega)|^2 \leq 2B, \quad \forall\, \omega \in \mathbb{R}^2 \setminus \{(0,0)\} \tag{82}$$

When the above semi-discrete frame conditions holds, the dyadic directional wavelet transform is a map $W : \mathbf{L}^2(\mathbb{R}^2) \to \boldsymbol{\ell}^2(\mathbf{L}^2(\mathbb{R}^2))$ defined as:

$$Wf = \{f * \bar{\psi}_{j,\alpha}(u) : j \in \mathbb{Z},\, \alpha \in \Theta,\, u \in \mathbb{R}^2\}, \quad \bar{\psi}_{j,\alpha}(x) = \psi^*_{j,\alpha}(-x)$$

A wavelet $\psi_{u,j,\alpha}(x) = \psi_{j,\alpha}(x-u)$ has support dilated by 2^j, located in a neighborhood of u and oscillates in the direction $\alpha + \pi/2$. If $f(x)$ is constant over the support of $\psi_{j,\alpha}(x-u)$ along lines of direction $\alpha + \pi/2$, then $f * \bar{\psi}_{j,\alpha}(u) = 0$ because of its directional vanishing moments. In particular, the wavelet coefficient vanishes in the neighborhood of an edge having a tangent in the direction of $\alpha + \pi/2$. If the edge angle deviates from $\alpha + \pi/2$, then it produces large amplitude coefficients, with a maximum typically when the edge has direction α. Figure 36 illustrates the idea.

Two dimensional Gabor wavelets are directional wavelets generated from a single complex valued mother wavelet. The mother wavelet is defined as:

$$\psi(x) = g_\sigma(x) e^{i\xi\cdot x}$$

where $\xi = (\xi_1, \xi_2) \in \mathbb{R}^2$ is the central frequency of ψ and $g(x)$ is a Gaussian, which we take as

$$g_\sigma(x) = \frac{1}{2\pi\sigma^2} e^{-|x|^2/2\sigma^2}$$

The mother Gabor wavelet oscillates along the angle $\arccos(\xi_1/|\xi|)$, and thus in the more general language of directional wavelets it has angle $\alpha = \arccos(\xi_1/|\xi|) - \pi/2$. The Fourier transform of ψ is

$$\widehat{\psi}(\omega) = e^{-\sigma^2|\omega-\xi|^2/2}$$

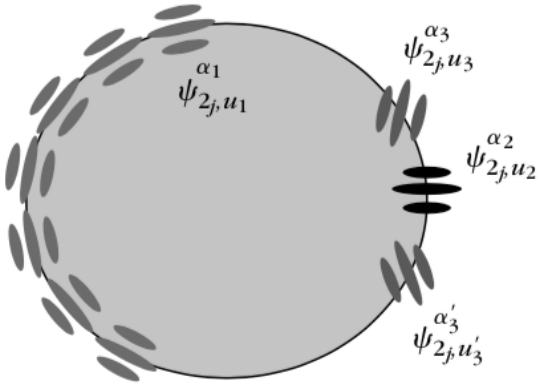

Figure 36: A cartoon image of a disk, with a regular edge. When the wavelet direction α is orthogonal to the tangent of the edge, the coefficients vanish as indicated by the black wavelet response (α_2). When the wavelet direction α aligns with the tangent of the curve (as on the left with α_1), the wavelet coefficients have large amplitude. When the tangent of the curve is not aligned with the wavelet, but is not orthogonal either (as in α_3 and α'_3), wavelet coefficients may have non-negligible amplitude but generally not as large as the α_1 coefficients.

For appropriate choices of $\sigma \in \mathbb{R}$ and $\xi \in \mathbb{R}^2$ the wavelet has nearly zero average and is almost analytic. The generators of a Gabor wavelet semi-discrete frame are obtained from dilations and rotations of the mother wavelet:

$$\psi_{j,\theta}(x) = 2^{-2j}\psi(2^{-j}R_\theta^{-1}x), \quad j \in \mathbb{Z},\, \theta \in \Theta$$

where R_θ is the two dimensional rotation matrix by the angle θ,

$$R_\theta = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$$

A simple computation shows that

$$\psi_{j,\theta}(x) = g_{2^j\sigma}(x)e^{i2^{-j}R_\theta\xi\cdot x}$$

Thus $\psi_{j,\theta}$ changes the essential support of ψ from a ball of radius σ to a ball of radius $2^j\sigma$, the direction of oscillation is rotated by θ radians, and the magnitude of the frequency of this oscillation is now $2^{-j}|\xi|$. In frequency we have:

$$\widehat{\psi}_{j,\theta}(\omega) = \widehat{\psi}(2^jR_\theta^{-1}\omega) = e^{-(2^j\sigma)^2|\omega-2^{-j}R_\theta\xi|^2/2}$$

Thus the essential support of $\widehat{\psi}_{j,\theta}(\omega)$ is a ball of radius $(2^j\sigma)^{-1}$ centered at $2^{-j}R_\theta\xi$. These frequency supports will cover the upper half plane for appropriate choices of K (the number

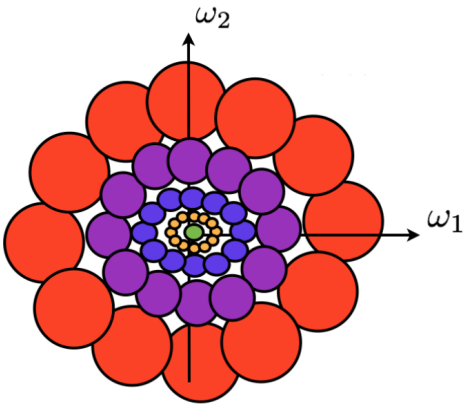

Figure 37: Frequency supports of the two dimensional oriented Gabor wavelets. Scales 2^j are indicated with different colors. The central green ball corresponds to the frequency support of a two dimensional scaling function.

of angles θ) and ξ; reflections will then cover the lower half plane, and so (82) will be satisfied. Figure 37 illustrates this frequency covering.

As in one dimension, Gabor wavelets can be amended to have precisely zero average:

$$\psi(x) = g_\sigma(x)(e^{i\xi\cdot x} - C), \quad C \text{ chosen so that } \int_{\mathbb{R}^2} \psi(x)\,dx = 0$$

which gives a directional Morlet wavelet. Figure 38 plots the real part of a Morlet wavelet at different scales and orientations; Figure 39 plots the imaginary part; and Figure 40 plots their Fourier transforms.

A dyadic Gabor/Morlet wavelet transform computes:

$$Wf = \{f * \psi_{j,\theta}(u) : j \in \mathbb{Z},\, \theta \in \Theta,\, u \in \mathbb{R}^2\}$$

Figure 41 shows the result of computing the dyadic Gabor wavelet transform of an image consisting of one texture embedded in another texture. The middle texture is relatively smooth along vertical lines, but has significant variations in the horizontal direction. It results that a Gabor wavelet transform with two directional angles $\theta = 0$ and $\theta = \pi/2$ will have large magnitude responses for $\theta = 0$, and negligible response for $\theta = \pi/2$. The outside texture, on the other hand, has the most variation along the angle $\alpha = \pi/4$. The Gabor wavelet coefficients along the directions $\theta = 0, \pi/2$ should be negligible for this outside texture, and indeed they are.

Figure 42 computes the Morlet wavelet transform of a black and white image of a butterfly, where the directional edge detection properties of the transform are exhibited, particularly in the wing of the butterfly.

As a point of comparison, we can define a non-directional wavelet $\psi(x)$ as the Laplacian of a Gaussian:

$$\psi(x) = -(\Delta g)(x)\,, \quad g(x) = \frac{1}{2\pi\sigma^2} e^{-|x|^2/2\sigma^2}$$

Figure 38: Real part of Morlet wavelets. Increasing scale left to right, and increasing angle in $[0, \pi)$ from top to bottom. Green is positive, pink is negative, and white is zero.

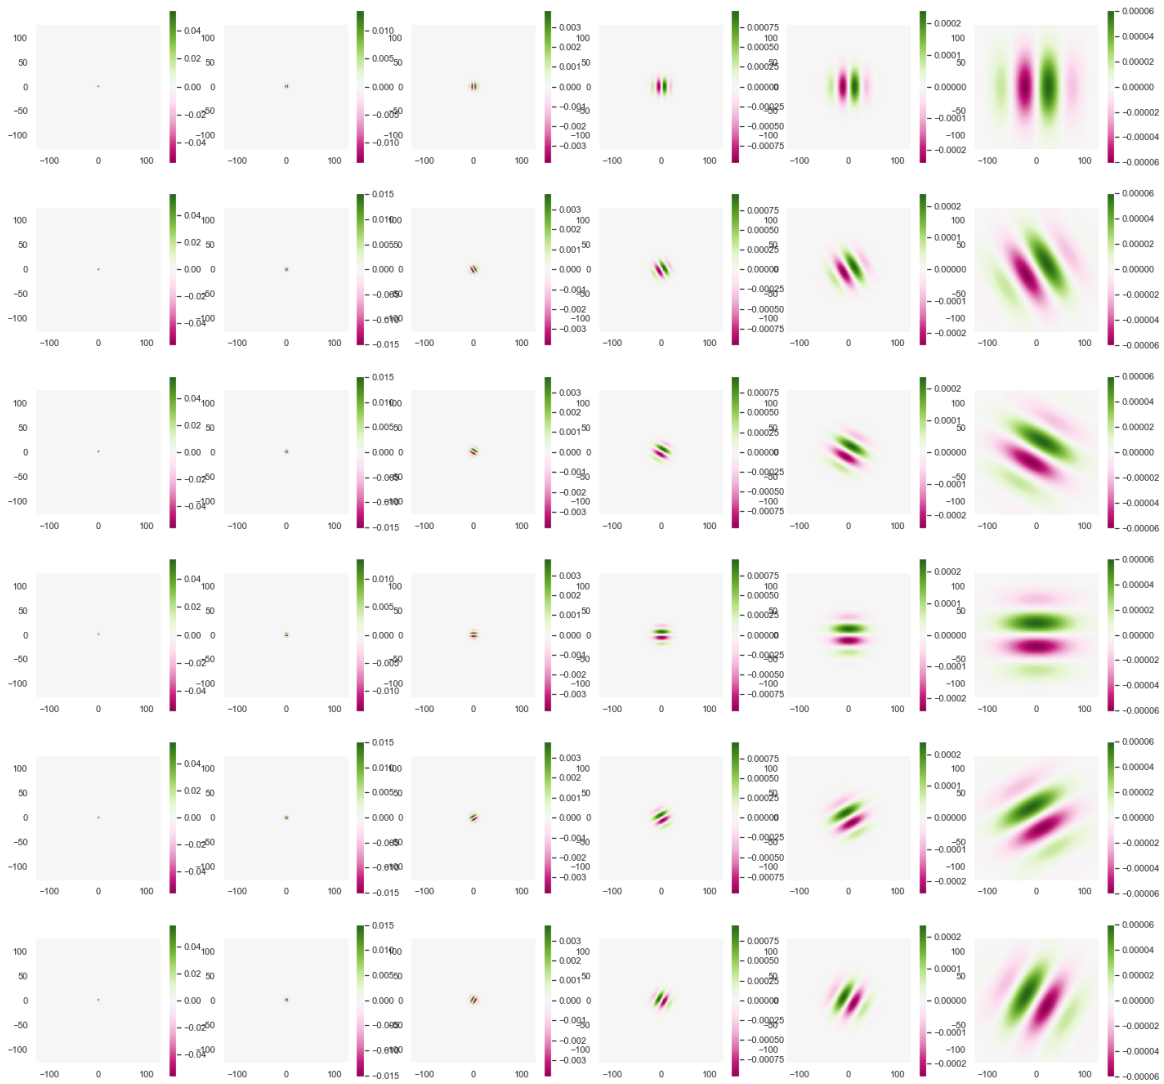

Figure 39: Imaginary part of Morlet wavelets. Increasing scale left to right, and increasing angle in $[0, \pi)$ from top to bottom. Green is positive, pink is negative, and white is zero.

Figure 40: Fourier transforms of Morlet wavelets. Increasing scale left to right, and increasing angle in $[0, \pi)$ from top to bottom. Green is positive, pink is negative, and white is zero.

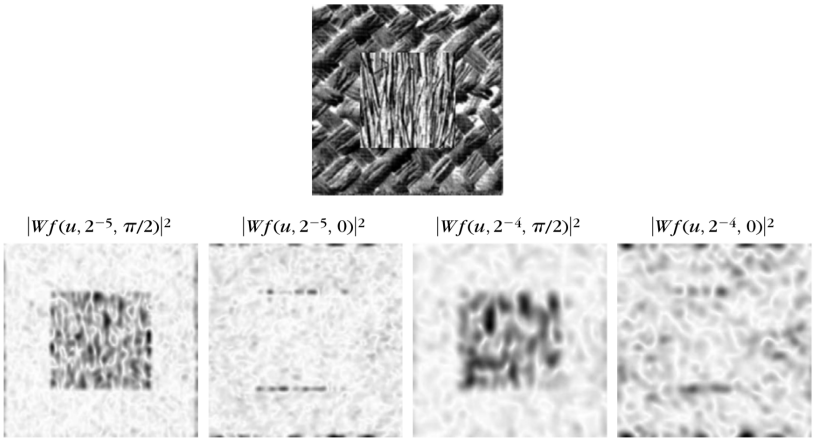

Figure 41: Top: Image of one texture embedded in another texture. Bottom: The magnitude of the Gabor wavelet transform $|Wf(u, j, \theta)| = |f * \psi_{j,\theta}(u)|$ for $j = -4, -5$ and $\theta = 0, \pi/2$. Images taken from Figure 5.9 of [1].

Note we have

$$\widehat{\psi}(\omega) = |\omega|^2 \widehat{g}(\omega) = |\omega|^2 e^{-\sigma^2 |\omega|^2/2}$$

and so ψ is indeed a wavelet, and is radially symmetric, meaning it has no directionality. The wavelet transform for this wavelet computes:

$$Wf = \{f * \psi_j(u) : j \in \mathbb{Z}\,,\, u \in \mathbb{R}^2\}$$

Even though this wavelet has no directionality, it can still pick up edges at small scales and meso- and macroscopic patterns at larger scales, agnostic of the direction. See Figure 43 for pictures and more details.

Exercise 74. Read Section 5.5.1 of *A Wavelet Tour of Signal Processing.*

Exercise 75. Read Section 5.5.2 of *A Wavelet Tour of Signal Processing.*

6.4 Multiscale Edge Detection

Section 6.3 of A Wavelet Tour of Signal Processing.

6.4.1 Wavelet Maxima for Images

Section 6.3.1 of A Wavelet Tour of Signal Processing.

Taken directly from [1]: *Image edges are often important for pattern recognition. This is clearly illustrated by our visual ability to recognize an object from a drawing that gives a*

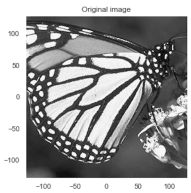

(a) Butterfly

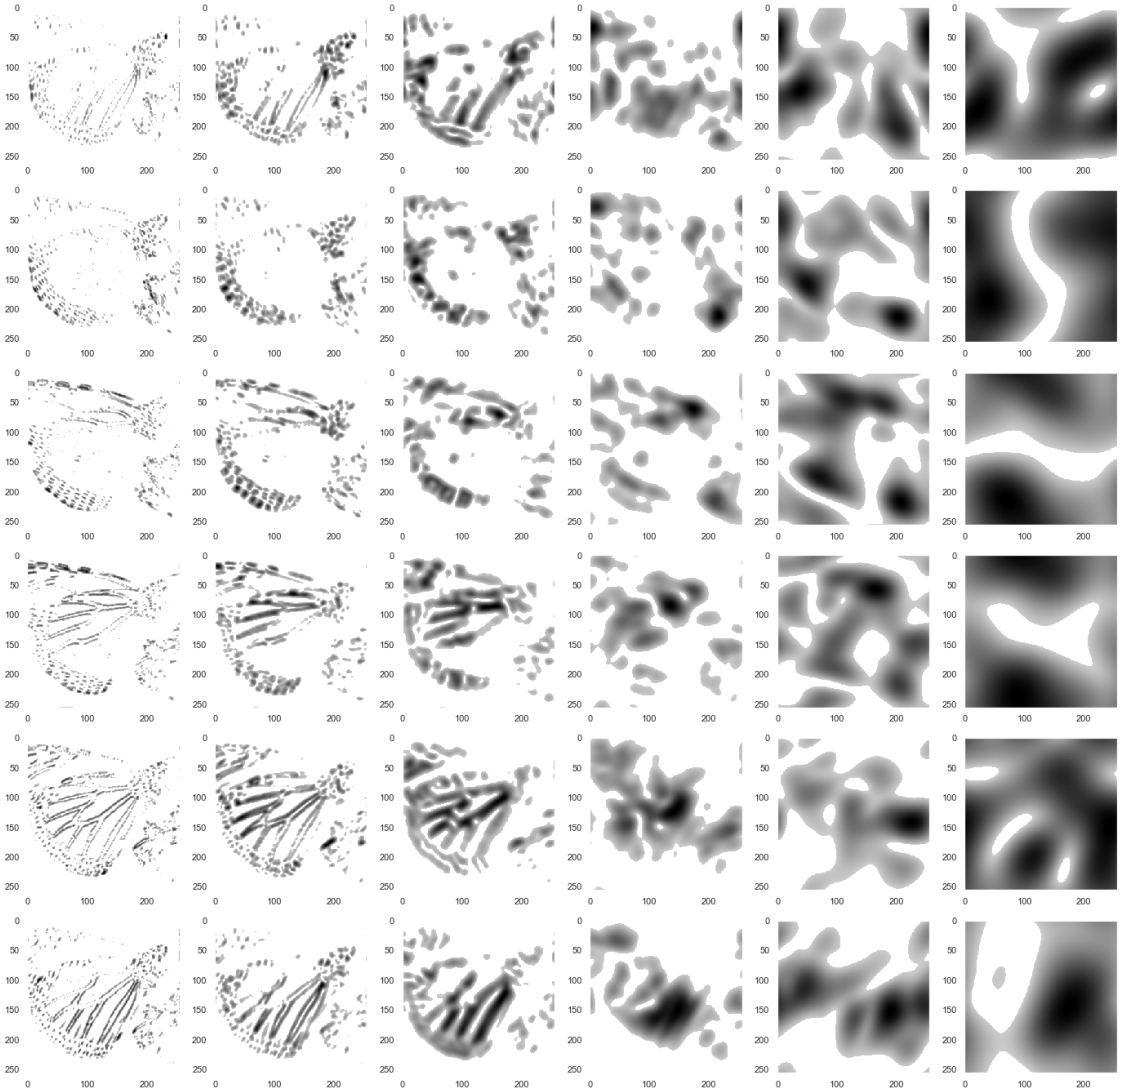

(b) Modulus of Morlet wavelet coefficients, $|f * \psi_{j,\theta}(u)|$.

Figure 42: (a) Black and white image of a butterfly. (b) Modulus of the Morlet wavelet coefficients, with increasing scale from left to right and increasing angle from top to bottom. Small wavelet coefficients have been set to zero to better illustrate the large amplitude coefficients.

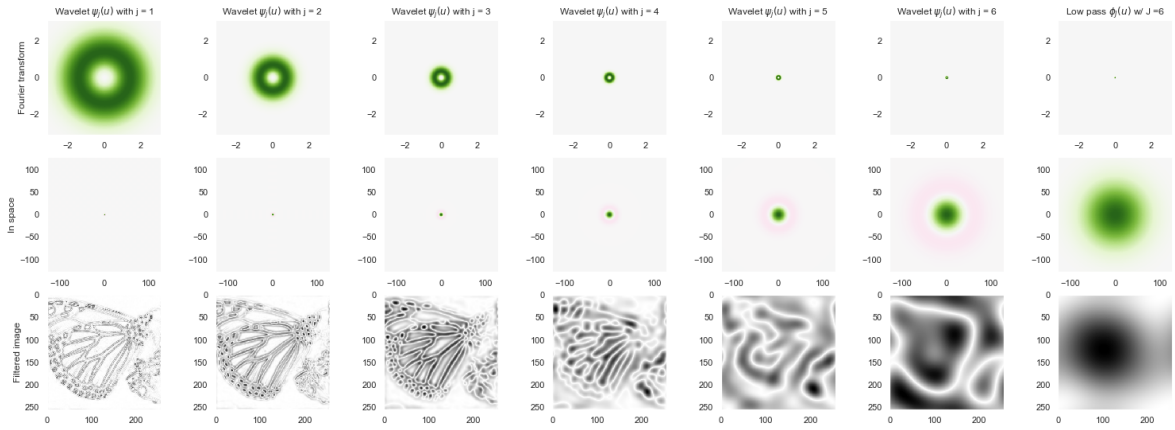

Figure 43: Non-directional wavelet transform based upon $\psi(x) = -(\Delta g)(x)$, for $g(x)$ a Gaussian. Scale increases from left to right; green is positive, pink is negative, white is zero in the first two rows. In the last row, grey/black is positive, white is zero. Top row: The Fourier transforms $\widehat{\psi}_j(\omega)$, which are like donuts. Middle row: The wavelets $\psi_j(x)$, which oscillate radially out from the origin. Bottom row: The absolute value of the wavelet coefficients, $|f * \psi_j(u)|$, for the butterfly image from Figure 42(a). Notice at small scales the wavelet still detects edges, but agnostic to the direction.

rough outline of contours. But, what is an edge? It could be defined as points where the image intensity has sharp transitions. A closer look shows that this definition is often not satisfactory. Image textures do have sharp intensity variations that are often not considered as edges. When looking at a brick wall, we may decide that the edges are the contours of the wall whereas the bricks define a texture. Alternatively, we may include the contours of each brick in the set of edges and consider the irregular surface of each brick as a texture. The discrimination of edges versus textures depends on the scale of analysis.

Figure 44 displays an image of bricks form the Brodatz image database.

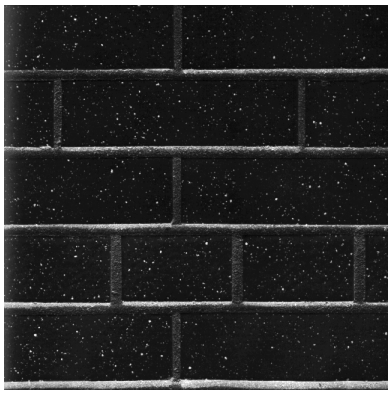

Figure 44: Bricks

Let $f : \mathbb{R}^2 \to \mathbb{R}$ be a two dimensional signal, such as an image. A Canny edge detection algorithm detects points of sharp variation in f by calculating the norm of the gradient vector of f. Recall the gradient of f is:

$$\nabla f(x) = \left(\frac{\partial f}{\partial x_1}(x), \frac{\partial f}{\partial x_2}(x) \right)$$

Let $\vec{n} = (\cos\beta, \sin\beta)$ be a unit vector in $\mathbb{R}^2$. The directional derivative of f in the direction $\vec{n}$ is defined as:

$$\frac{\partial f}{\partial \vec{n}}(x) = \nabla f(x) \cdot \vec{n} = \frac{\partial f}{\partial x_1}(x) \cos\beta + \frac{\partial f}{\partial x_2}(x) \sin\beta$$

The absolute value of $\partial f / \partial \vec{n}$ is maximum if $\vec{n}$ is colinear to ∇f. Since $\partial f / \partial \vec{n}(x)$ will be maximum in the direction of most change around x, this shows that the gradient $\nabla f(x)$ points in the direction of the most rapid increase around the point $x \in \mathbb{R}^2$; see Figure 45 for an illustration. The magnitude of the variation at x (or edge strength) is given by the norm of the gradient:

$$|\nabla f(x)| = \sqrt{\left(\frac{\partial f}{\partial x_1}(x) \right)^2 + \left(\frac{\partial f}{\partial x_2}(x) \right)^2}$$

Additionally the direction of the gradient is given by:

$$Af(x) = \begin{cases} \tan^{-1}\left(\frac{\partial f/\partial x_2(x)}{\partial f/\partial x_1(x)} \right) & \frac{\partial f}{\partial x_1}(x) \geq 0 \\ \pi + \tan^{-1}\left(\frac{\partial f/\partial x_2(x)}{\partial f/\partial x_1(x)} \right) & \frac{\partial f}{\partial x_1}(x) < 0 \end{cases} \tag{83}$$

A point $y \in \mathbb{R}^2$ is defined as an edge if $|\nabla f(x)|$ is locally maximum at $x = y$. Figure 46 shows the performance of the Canny edge detector (a refined version of it) on a color image.

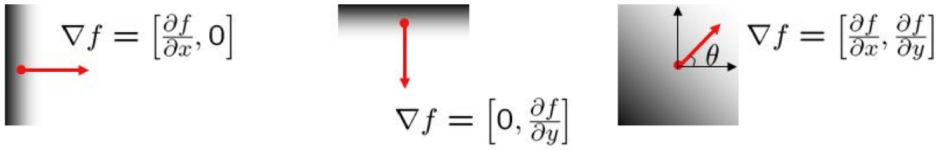

Figure 45: Illustration of the gradient vector for simple images.

A multiscale version of the Canny edge detector is implemented by smoothing the image with a convolution kernel $\theta(x)$ that is dilated at different scales. This is computed with two wavelets that are partial derivatives of θ:

$$\psi_1 = -\frac{\partial \theta}{\partial x_1}, \qquad \psi_2 = -\frac{\partial \theta}{\partial x_2}$$

These wavelets are dilated at the scales 2^j for $j \in \mathbb{Z}$:

$$\psi_{j,k}(x) = 2^{-2j}\psi_k(2^{-j}x)$$

and the dyadic wavelet transform computes:

$$Wf = \{f * \bar{\psi}_{j,k}(u) : j \in \mathbb{Z},\ k = 1, 2,\ u \in \mathbb{R}^2\}, \quad \bar{\psi}_{j,k}(x) = \psi_{j,k}(-x)$$

We write $Wf(u, j)$ as

$$Wf(u, j) = \begin{pmatrix} f * \bar{\psi}_{j,1}(u) \\ f * \bar{\psi}_{j,2}(u) \end{pmatrix}$$

and we think of it as two-dimensional vector valued function, $Wf(u, j) \in \mathbb{R}^2$. The wavelets $\psi_{j,1}$ measure variations in the horizontal direction at the scale 2^j, while the wavelets $\psi_{j,2}$ measure variations in the vertical direction at the scale 2^j.

Denote

$$\theta_j(x) = 2^{-2j}\theta(2^{-j}x), \qquad \bar{\theta}_j(x) = \theta_j(-x)$$

Figure 46: The result of the Canny edge detector applied to the image on the left. Taken from https://en.wikipedia.org/wiki/Canny_edge_detector

The dyadic wavelets can be written as:

$$\bar{\psi}_{j,k} = 2^j \frac{\partial \bar{\theta}_j}{\partial x_k}$$

It follows that

$$f * \bar{\psi}_{j,k}(u) = 2^j \frac{\partial}{\partial u_k}(f * \bar{\theta}_j)(u)$$

and so

$$Wf(u,j) = \begin{pmatrix} f * \bar{\psi}_{j,1}(u) \\ f * \bar{\psi}_{j,2}(u) \end{pmatrix} = 2^j \nabla (f * \bar{\theta}_j)(u)$$

Therefore $Wf(u,j)$ is proportional (up to a factor 2^j) to the gradient of the smooth version of f at the scale 2^j. The norm of $Wf(u,j)$, defined as:

$$|Wf(u,j)| = \sqrt{|f * \bar{\psi}_{j,1}(u)|^2 + |f * \bar{\psi}_{j,2}(u)|^2}$$

is thus proportional to $|\nabla(f * \bar{\theta}_j)(u)|$ as well. We can also compute the angle $AWf(u,j)$ of $Wf(u,j)$ using the definition (83).

The unit vector

$$\vec{n}_j(u) = (\cos AWf(u,j), \sin AWf(u,j))$$

is colinear to $\nabla(f * \bar{\theta}_j)(u)$. An edge point at the scale 2^j is a point $v \in \mathbb{R}^2$ such that $|Wf(u,j)|$ is locally maximum at at $u = v$. These points are two dimensional wavelet modulus maxima. Individual wavelet modulus maxima are chained together to form a maxima curve that follows an edge.

Figure 47 computes the Canny dyadic wavelet transform of image of a disc. The wavelet modulus maxima curves are along the boundary of the disc.

As in 1D the decay of the 2D Canny dyadic wavelet coefficients depends upon the local regularity of f. Let $0 \leq \alpha \leq 1$ denote the Lipschitz regularity. A function $f : \mathbb{R}^2 \to \mathbb{R}$ is Lipschitz α at $v \in \mathbb{R}^2$ if there exists $K > 0$ such that for all $x \in \mathbb{R}^2$,

$$|f(x) - f(v)| \leq K|x - v|^\alpha$$

where $|x|$ is the norm of $x \in \mathbb{R}^2$. As in one dimension, the local Lipschitz regularity of f is related to the asymptotic decay of $Wf(u,j)$. Indeed, this regularity is controlled by $|Wf(u,j)|$. Let $\Omega \subset \mathbb{R}^2$ be a bounded domain. One can prove that f is uniformly Lipschitz α inside Ω if and only if there exists $A > 0$ such that

$$|Wf(u,j)| \leq A 2^{j\alpha}, \quad \forall\, u \in \Omega$$

Also analogously to the 1D setting, we can synthesize high fidelity approximations to images using only their wavelet modulus maxima. Similarly to before, let Λ denote the set

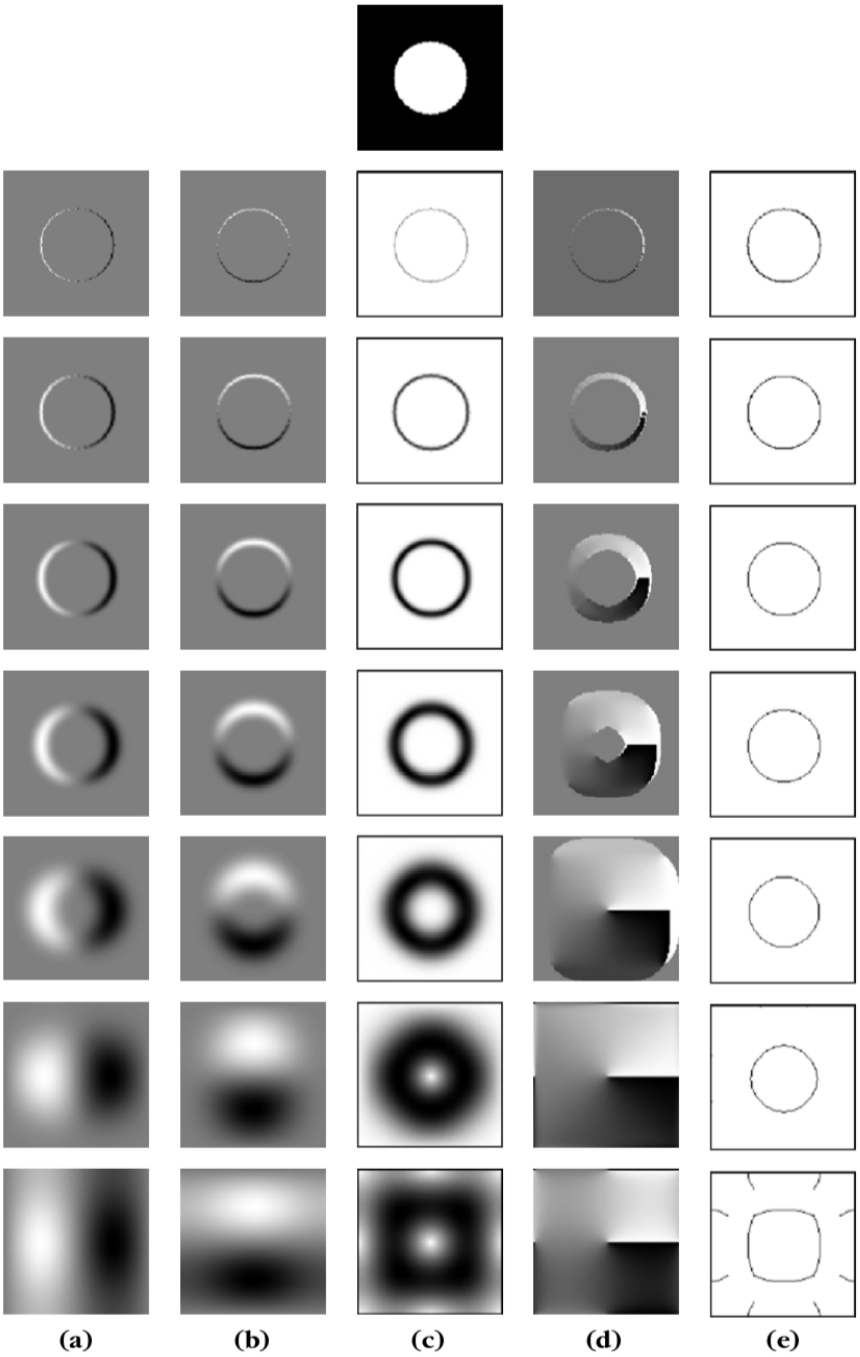

Figure 47: Top image: The disc. (a) Horizontal wavelet transform: $f * \bar{\psi}_{j,1}(u)$ for $-6 \leq j \leq 0$. (b) Vertical wavelet transform: $f * \bar{\psi}_{j,2}(u)$ for $-6 \leq j \leq 0$. (c) Norm of wavelet coefficients $|Wf(u,j)|$. (d) Angles $AWf(u,j)$ where $|Wf(u,j)| \neq 0$. (e) The wavelet modulus maxima curves.

of wavelet modulus maxima of f using the Canny dyadic wavelet transform. The synthesized image is computed using the dual synthesis:

$$f_\Lambda = \sum_{(u,j)\in\Lambda} \sum_{k=1,2} \langle f, \psi_{u,j,k}\rangle \widetilde{\psi}_{u,j,k}$$

One can also de-noise images by thresholding wavelet modulus maxima. Suppose that we have only a noisy version of f, modeled as

$$\widetilde{f}(x) = f(x) + \varepsilon(x), \quad \varepsilon(x) \sim \mathcal{N}(0, \sigma^2)$$

where $\varepsilon(x)$ is sampled independently and identically from the normal distribution; ε is referred to as white noise. The wavelet coefficients of $\widetilde{f}$ are:

$$W\widetilde{f} = Wf + W\varepsilon$$

At the large scales the averaging by $\bar{\theta}_j$ kills much of $W\varepsilon$ since ε has zero average. However, at small scales, the wavelets respond against ε and $W\varepsilon$ masks Wf. The image can be de-noised by thresholding the wavelet coefficients $W\widetilde{f}$, which results in a "cartoon" version of the original image if the variance σ^2 of the noise is not too large; Figure 48 illustrates the idea.

Exercise 76. Read Section 6.3 of *A Wavelet Tour of Signal Processing*.

Exercise 77. OPTIONAL Using your code from previous exercises compute the dyadic wavelet transform of the signal in Figure 13. Compute the wavelet modulus maxima as well. Implement a dual synthesis projection (however you like) and compute f_Λ, i.e., the signal synthesized from only the wavelet modulus maxima coefficients. Threshold the wavelet modulus maxima coefficients and synthesize a signal only from the largest ones. Turn in plots of the wavelet coefficients, the wavelet modulus maxima, and the synthesized signals. Explain your results.

Exercise 78. OPTIONAL Let $K \in \mathbb{N}$ with $K \geq 2$ and define $\vec{n}_k = (\cos(2\pi k/K), \sin(2\pi k/K)) \in \mathbb{R}^2$.

(a) Prove that $\mathcal{D} = \{\vec{n}_k : 0 \leq k < K\}$ is a tight frame of K vectors in $\mathbb{R}^2$ and that for any $\omega \in \mathbb{R}^2$, it satisfies

$$\sum_{k=0}^{K-1} |\omega \cdot \vec{n}_k|^2 = \frac{K|\omega|^2}{2}$$

(b) Let

$$\psi_k = \frac{\partial \theta}{\partial \vec{n}_k}$$

be the directional derivative of $\theta(x)$ in the direction $\vec{n}_k$. Define dilations of ψ_k as:

$$\psi_{j,k}(x) = 2^{-2j}\psi_k(2^{-j}x), \quad j \in \mathbb{Z}$$

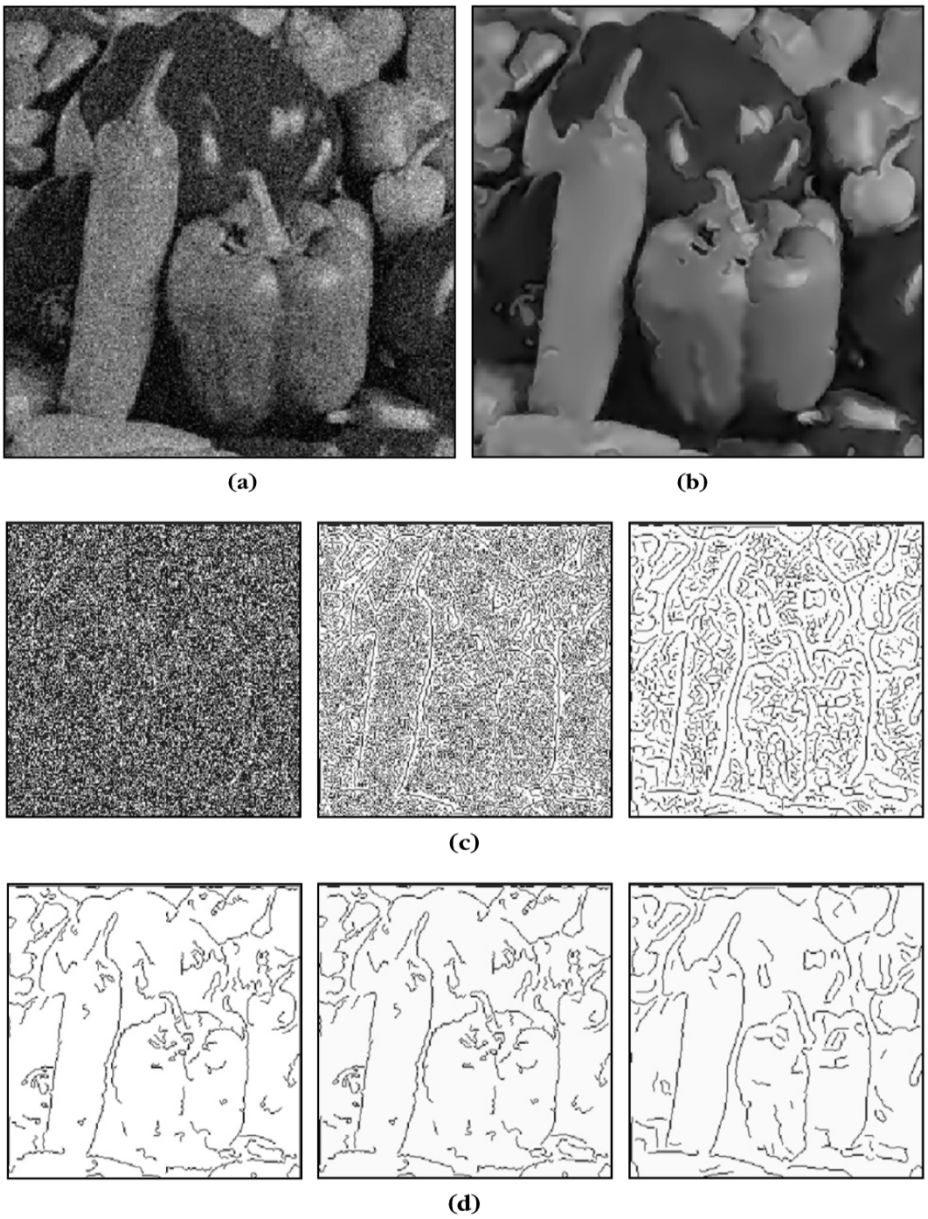

Figure 48: (a) Noisy peppers image. (b) Restored peppers image from the thresholding the maxima curves shown in (d). (c) The wavelet modulus maxima points of the noisy image for scales $-7 \leq j \leq -5$. (d) The thresholded wavelet modulus maxima.

If $\theta(x)$ is rotationally invariant (i.e., $\theta(x)$ depends only on $|x|$), then prove that $\mathcal{D} = \{\psi_{j,k} : j \in \mathbb{Z}, 0 \leq k < K\}$ are the generators of translation invariant semi-discrete frame if and only if

$$\frac{2A}{K} \leq \sum_{j \in \mathbb{Z}} 2^{2j} |\omega|^2 |\widehat{\theta}(2^j \omega)|^2 \leq \frac{2B}{K}, \quad \text{a.e. } \omega \in \mathbb{R}^2$$

Exercise 79. **OPTIONAL** In this problem you will compute the wavelet transform of two dimensional textures.

(a) Implement a dyadic 2D Morlet wavelet transform. Visualize your wavelets and their Fourier transforms, as in the plots posted in the #in-class channel of the course Slack, and turn these visualizations in.

(b) Take an image of your choice, and compute the dyadic Morlet wavelet transform of the image. Turn in plots of the real and imaginary parts of the wavelet coefficients for each (j, θ), and the modulus of the wavelet coefficients for each (j, θ). Do you see the directional responses at different scales?

References

[1] Stéphane Mallat. *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*. Academic Press, 3rd edition, 2008.

[2] Elias M. Stein and Rami Shakarchi. *Fourier Analysis: An Introduction*. Princeton Lectures in Analysis. Princeton University Press, 2003.

[3] John J. Benedetto and Matthew Dellatorre. Uncertainty principles and weighted norm inequalities. *Contemporary Mathematics*, 693:55–78, 2017.

[4] Yves Meyer. *Wavelets and Operators*, volume 1. Cambridge University Press, 1993.

[5] Karlheinz Gröchenig. *Foundations of Time Frequency Analysis*. Springer Birkhäuser, 2001.

[6] Steven Shreve. *Stochastic Calculus for Finance II*. Springer-Verlag New York, 2004.

[7] Hermine Biermé. Introduction to random fields and scale invariance. hal-01493834, 2018.

[8] Georgiy Shevchenko. Fractional Brownian motion in a nutshell. arXiv:1406.1956, 2014.