

Discrete Directional Gabor Systems for Image Processing

James M. Murphy

Department of Mathematics
Johns Hopkins University

October 6, 2016

Overview

- First, we will present a mathematical construction of discrete directional Gabor frames, introducing sufficient conditions on the window function and sampling set Λ_ω .
- Second, these frames will be evaluated against state-of-the-art frames for a variety of images.
- This is joint work with Wojciech Czaja (UMD), Benjamin Manning (UMD), and Kevin Stubbs (Duke).

Mathematical Motivation

- Given a signal f , one wishes to extract its most important features for analysis.
- If f represents an image, some important analysis tasks include segmentation, registration, classification, superresolution, denoising, and compression.
- One class of methods that extract useful features from f are methods of *harmonic analysis*, in which one represents f in an illuminating way.
- Classical examples:
 - 1 Fourier analysis
 - 2 Wavelet analysis
 - 3 Time-frequency (Gabor) analysis

Frame Theory

- The notion of *frame* encodes the idea of feature extraction and signal representation in a way amenable to mathematical analysis.

Definition

Let \mathcal{H} be a Hilbert space. A discrete set $\{\phi_i\}_{i \in I} \subset \mathcal{H}$ is called a (*discrete*) *frame* for \mathcal{H} if there exist constants $0 < A \leq B < \infty$ such that:

$$\forall f \in \mathcal{H}, A\|f\|_{\mathcal{H}}^2 \leq \sum_{i \in I} |\langle f, \phi_i \rangle_{\mathcal{H}}|^2 \leq B\|f\|_{\mathcal{H}}^2.$$

The optimal choices of A, B are the *frame bounds*

- For the analysis of 2D images, $\mathcal{H} = L^2([0, 1]^2)$. We will focus on this case unless otherwise stated, though our theoretical results hold for $L^2([0, 1]^d)$.

Classical Frames

- Many important representation systems fit into this definition.
- **Fourier frames:** $\{e^{i\lambda x}\}_{\lambda \in \Lambda}$. For $\Lambda = (2\pi m, 2\pi n) | m, n \in \mathbb{Z}$, this is Fourier series.
- **Wavelet frames:** $\{\psi(A^k(x - (n, m)))\}_{k, n, m \in \mathbb{Z}}$, for some wavelet function ψ and dilation matrix A .
- **Gabor frames:** $\{g(x - \omega)e^{i\lambda x}\}_{\omega \in \Omega, \lambda \in \Lambda}$. This can be considered a *windowed* Fourier frame.

Frames and Image Processing

- Frames offer a principled system for representing data: given a frame $\{\phi_i\}_{i \in I}$, write a signal f as

$$f = \sum_{i \in I} \langle f, \phi_i \rangle \tilde{\phi}_i.$$

for a set of *dual frame elements* $\{\tilde{\phi}_i\}$.

- Different frames are known to provide efficient representations of different datasets.
- If a scientific user knows something about the structure of their data a priori, then one can use a frame well-suited to that class of data.
- As a computational method, once a frame has been chosen, fast algorithms for implementation are often achievable.

Simple Frame Compression

- Compression is a standard application of frames. Writing

$$f = \sum_{i \in I} \langle f, \phi_i \rangle \tilde{\phi}_i.$$

for a set of *dual frame elements* $\{\tilde{\phi}_i\}$, we can reconstruct f from its *frame coefficients* $\{\langle f, \phi_i \rangle\}$.

- **Question:** Do we need *all* the coefficients to reconstruct f well?
- **Answer:** No, if f is a signal well-suited to the frame $\{\phi_i\}$; this can be made precise via the notion of sparsity in a frame. In this case, many coefficients ≈ 0 .

Fourier Compression

- A simple illustration of this occurs with Fourier series in one dimension.
- Let $f(t) = \sin(2\pi t)$. Instead of storing a discretized version of f , one can simply store a couple of Fourier coefficients:

$$\langle f, e^{-2\pi im} \rangle = \int_0^1 \sin(2\pi t) e^{2\pi it} dt = \begin{cases} \frac{i}{2}, & m = 1 \\ -\frac{i}{2}, & m = -1 \\ 0, & \text{else} \end{cases}$$

- This of course makes sense, since Fourier series represent signals as sums of $\sin(t)$ and $\cos(t)$ of various frequencies, and our signal $f(t)$ is a pure $\sin(t)$.
- We say that the signal f has a *sparse* Fourier representation.
- This idea extends to richer classes of signals for other frames.

Anisotropic Frames

- Over the past decade, several new frames were developed that proved very efficient for image processing tasks on images with *smooth edges*: shearlets, curvelets, etc.
- More precisely, for a certain class of data that models signals with smooth edges, shearlets and ridgelets are known to provide near-optimal sparsity.
- This image class is known as *cartoon-like* images in the literature.
- This allows them to compress and denoise these images very efficiently.

Frames for Textured Images

- While these frames are effective for images with edges, they are not known to perform well for *textured images*.
- Many interesting images fall into this regime, and thus there is a need to develop anisotropic frames well-suited for textures.
- Building on the knowledge that Gabor systems perform well empirically for many image processing tasks with isotropic textural images, we proposed to incorporate directionality into the construction.

Directional Gabor frame elements

Definition

For $g : \mathbb{R} \rightarrow \mathbb{C}$ and $m, t \in \mathbb{R}$, $u \in \mathcal{S}^{d-1} = \{x \in \mathbb{R}^d \mid x^1 + x^2 + \dots + x^d = 1\}$,

$$g_{m,t,u}(x) := e^{2\pi i m(u \cdot x)} g(u \cdot x - t), \quad m, t \in \mathbb{R}, \quad u \in \mathcal{S}^{d-1}, \quad x \in \mathbb{R}^d$$

- We seek a discrete system of the form

$$\{g_{m,t,u}\}_{(m,t,u) \in \Lambda},$$

for $L^2([0, 1]^d)$. We need to determine g, Λ that work.

- Of course, $d = 2$ is of particular interest for image processing.

Existence of Frames

- Let $\Gamma \subset \mathcal{S}^{d-1} \times \mathbb{R}$ be such that the mapping $\psi : \Gamma \rightarrow \mathbb{Z}^d$ given by $(u, m) \mapsto mu$ is a bijection. For any $\omega > 0$, let

$$\Lambda_\omega = \{(m, n, u) \mid (u, m) \in \Gamma, n \in \omega\mathbb{Z}\}.$$

Theorem

Let $g \in L^2(\mathbb{R})$ be compactly supported and such that

$$\forall k \in (\mathbb{Z}/\omega) \setminus \{0\}, \hat{g}(\gamma)\hat{g}(\gamma + k) = 0$$

almost everywhere. Furthermore, suppose it is not the case that \hat{g} is zero almost everywhere on the interval $[-1/4, 1/4]$. Then with Λ_ω as above, $\{g_{m,n,u}\}_{(m,n,u) \in \Lambda_\omega}$ is a discrete frame for the subspace of functions in $L^2(\mathbb{R}^d)$ with support contained in $[-1/2, 1/2]^d$.

- Note that $[-\frac{1}{2}, \frac{1}{2}]^d$ may be replaced with $[0, 1]^d$, and a frame can still be shown to exist, using the same methods.

Image Processing Applications

- We now consider two image processing applications for our frame: compression and denoising.
- To evaluate our method, we compare against several well-known isotropic and anisotropic frames: wavelets, shearlets, and curvelets.
- We consider images that are highly textural and also more natural ones with more edges and textures.

Test Images (1/2)




Figure: Texture images for numerical experiments. Top row, from left to right: straw, rocks, grass. Middle row, from left to right: cracked mud, bricks, bars. Bottom row, from left to right: fabric, grate, honeycomb.

Test Images (2/2)



Figure: Real images for numerical experiments, from left to right: mandrill, boat, Lena.

Compression Experiments

- To perform compression, we compute the discrete frame coefficients, then perform hard thresholding.
- That is, we set a certain percentage of the smallest coefficients to 0, then reconstruct by computing the inverse frame operator.
- For our experiments, four levels of compression are considered, parameterized by compression ratios of 10, 25, 50, 100. These correspond to 90%, 96%, 98%, 99% thresholding, respectively. We evaluate quality by computing the mean squared error (ℓ^2 error) between the compressed and original images.
- The higher the level of compression, the more significant our storage advantage. Conversely, higher levels of compression lead to higher reconstruction errors.

Compression Results

Method	Compression Level	Straw	Rocks	Grass	Cracked Mud	Bricks	Bars	Fabric	Grate	Honeycomb	Mandrill	Boat	Lena
DGS, nonredundant	90%	.1099	.0353	.1451	.1501	.0851	.0395	.1699	.1381	.0361	.1799	.1288	.1332
	96%	.1574	.0561	.2032	.1960	.1021	.0598	.2140	.1728	.0541	.2051	.1570	.1719
	98%	.1921	.0778	.2407	.2225	.1138	.0869	.2366	.1964	.0696	.2174	.1782	.2026
	99%	.2212	.1042	.2703	.2437	.1253	.1299	.2521	.2196	.0870	.2279	.2002	.2373
DGS, redundant	90%	.0453	.0158	.0597	.0633	.0412	.0170	.0692	.0640	.0148	.0906	.0619	.0642
	96%	.0859	.0292	.1107	.1169	.0693	.0324	.1309	.1113	.0288	.1464	.1043	.1098
	98%	.1176	.0410	.1525	.1538	.0865	.0457	.1695	.1416	.0408	.1766	.1310	.1401
	99%	.1509	.0563	.1934	.1855	.1005	.0630	.2018	.1688	.0546	.1985	.1541	.1688
Shearlab	90%	.1191	.0521	.1808	.1613	.0912	.0439	.1651	.1675	.1003	.1509	.1079	.1380
	96%	.1946	.0961	.2542	.2222	.1335	.1166	.2364	.2581	.2582	.2105	.1682	.2303
	98%	.2784	.1528	.3136	.2824	.2482	.2861	.3079	.3443	.3974	.2516	.2379	.3202
	99%	.6367	.5639	.6131	.6250	.6541	.6699	.6647	.6675	.6879	.5590	.5694	.5264
FFST	90%	.1292	.0820	.1926	.1742	.0979	.0467	.1864	.1925	.1115	.1769	.1300	.1607
	96%	.2393	.1760	.3025	.2600	.1849	.1918	.2778	.3104	.3589	.2767	.2392	.3354
	98%	.6037	.5776	.6061	.5874	.5786	.6274	.6085	.6209	.6655	.5910	.5230	.5771
	99%	.7860	.7668	.7884	.7873	.7961	.8015	.7991	.8164	.8413	.7679	.7250	.7438
Curvelab	90%	.1297	.0378	.1613	.1593	.1232	.1255	.2013	.2224	.2334	.1370	.0979	.1059
	96%	.1822	.0613	.2221	.2112	.1829	.1874	.2582	.2906	.3440	.1898	.1469	.1687
	98%	.3062	.1709	.3136	.3134	.3186	.3549	.3584	.3977	.4562	.2734	.2243	.2499
	99%	.6395	.5963	.6147	.6380	.6729	.6959	.6847	.7028	.7435	.5959	.5926	.5070
Wavelet	90%	.1569	.0407	.1607	.1504	.1084	.1148	.2019	.2546	.3033	.1528	.0937	.1008
	96%	.2008	.0642	.2137	.1991	.1642	.1804	.2461	.2967	.3534	.1926	.1380	.1609
	98%	.2261	.0859	.2482	.2277	.1957	.2107	.2677	.3155	.3728	.2136	.1682	.2043
	99%	.2469	.1134	.2771	.2521	.2168	.2331	.2824	.3275	.3838	.2320	.1962	.2473

Table: Compression errors.

Visual Evaluation of Compression




Figure: Compressed images for straw texture. The rows correspond to level of compression: from top to bottom, 90%, 96%, 98%, 99%. The columns correspond to method; from left to right: non-redundant directional Gabor frame, redundant directional Gabor frame, ShearLab, FFST, Curvelab, wavelet.

Image Denoising

- Frames are also useful denoising tools, motivated by the notion that large-valued frame coefficients values represent signal, and small-valued frame coefficients represent noise.
- One can then attempt to remove the noise while preserving the signal by reconstructing only with the largest frame coefficients; this is known as *hard-thresholding frame denoising*:

$$\tilde{\mathbf{f}} = \sum_{i \in \tilde{\mathcal{I}}} \langle \mathbf{f}, \phi_i \rangle \tilde{\phi}_i.$$

for the largest frame coefficients $\{\langle \mathbf{f}, \phi_i \rangle\}_{i \in \tilde{\mathcal{I}}}$.

- We consider experiments in which our images have been synthetically corrupted with Gaussian (white) noise, and we attempt to reconstruct the original image as accurately as possible. Quality of performance is measured in peak-signal-to-noise ratio (PSNR), which is basically rescaled mean-square error.

Denosing Results

Image	DGS, non-redundant	DGS redundant	SL PSNR	FFST PSNR	CL PSNR	Wavelet PSNR
Straw	16.1568	16.9645	15.9545	15.8711	15.5610	15.2016
Rocks	19.6596	19.6498	18.2772	15.9770	19.3774	19.0008
Grass	18.0498	18.1756	17.9382	17.4305	17.9655	17.9847
Cracked Mud	15.5964	16.1645	15.2524	15.0387	15.1937	15.2714
Bricks	21.8677	20.9070	19.7579	19.5595	18.4367	18.9953
Bars	21.6415	21.1399	19.0507	19.2957	17.6218	18.2207
Fabric	16.2813	16.5387	15.7118	15.6807	15.2957	15.5106
Grate	17.7373	17.8509	15.5762	15.6462	14.7761	14.9119
Honeycomb	19.9771	19.6670	14.6194	15.2500	12.8580	11.7106
Mandrill	16.5017	16.5743	16.4203	15.1830	16.5653	16.2576
Boat	17.9514	18.0641	17.6947	16.4199	18.2324	18.0268
Lena	17.3064	17.6526	16.5380	15.3970	17.1781	16.7578

Table: Best-case thresholding PSNR for denosing experiments.

Visual Evaluation of Denoising









Figure: Denoising results for straw. First row, left to right: original, noisy, denoised with non-redundant directional Gabor system, denoised with redundant directional Gabor system. Second row, left to right: denoised with shearlab shearlets, denoised with FFST shearlets, denoised with curvelab curvelets, denoised with wavelets.





Conclusions and Future: Theoretical Guarantees, Fast Algorithms, Applications

- Incorporating directionality into Gabor systems generates frames that offer strong performance for compression and denoising.
- This is particularly so for textural images, but for images with both textures and edges, results are more mixed.
- Developing a theory for why directional Gabor systems perform well numerically for textural images is the outstanding mathematical question.
- The algorithm we deploy for our frame requires a numerical inversion of a matrix to compute the frame reconstruction, which is slower than many available frame representation packages; developing a fast algorithm will greatly improve our method's applicability.
- We have tested our method on standard image processing images but finding a rich class of real textured but noisy images would extend the applicability of our method.





References I

-  A. Baudes, B. Coll, and J.M. Morel.
A review of image denoising algorithms, with a new one.
Multiscale Model. and Simul., 4(2):490–530, 2005.
-  E.J. Candès and D.L. Donoho.
Continuous curvelet transform: II. Discretization and frames.
Applied and Computational Harmonic Analysis, 19(2):198–222, 2005.
-  E. Candès, L. Demanet, D. Donoho, and L. Ying.
Fast discrete curvelet transforms.
Multiscale Modeling & Simulation, 5(3):861–899, 2006.
-  W. Czaja, B. Manning, J.M. Murphy, and K. Stubbs.
Discrete directional Gabor frames.
To appear in Applied and Computational Harmonic Analysis, 2016*.
-  D.L. Donoho and J.M. Johnstone.
Ideal spatial adaptation by wavelet shrinkage.
Biometrika, 81(3):425–455, 1994.

References II

-  R.J. Duffin and A.C. Schaeffer.
A class of nonharmonic Fourier series.
Transactions of the American Mathematical Society, pages 341–366, 1952.
-  M.N. Do and M. Vetterli.
The finite ridgelet transform for image representation.
IEEE Transactions on Image Processing, 12(1), 2003.
-  D. Gabor.
Theory of communication. Part 1: The analysis of information.
Journal of the Institution of Electrical Engineers -Part III: Radio and Communication Engineering, 93(26):429–441, 1946.
-  K. Gröchening.
Foundations of Time-Frequency Analysis.
Birkhäuser, 2001.

References III

-  L. Grafakos and C. Sansing.
Gabor frames and directional time-frequency analysis.
Applied and Computational Harmonic Analysis, 25(1):47–67, 2008.
-  S. Häuser and G. Steidl.
Fast finite shearlet transform.
arXiv preprint arXiv:1202.1773, 2012.
-  G. Kutyniok and D. Labate, editors.
Multiscale analysis for multivariate data.
Springer Science & Business Media, 2012.
-  G. Kutyniok, M. Shahram, and X. Zhuang.
Shearlab: A rational design of a digital parabolic scaling algorithm.
SIAM Journal on Imaging Sciences, 5(4):1291–1332, 2012.