Gabor Filter Visualization

V. Shiv Naga Prasad*
University of Maryland
shivnaga@cs.umd.edu

Justin Domke
University of Maryland
domke@cs.umd.edu

ABSTRACT

We present a system for the visualization of an important signal processing technique- a Gabor Filter bank's response to an image. To do this, one must overcome the problem that no multi-dimensional space can be shown in a a single, static graph. We use an interactive widget to change the visible range of the projected dimensions, and additional graphics which summarize the responses in the projected dimensions. Thus, though we view this four-dimensional space through 2-dimensional projections, we allow the user to understand all dimensions, not just the plane of projection. We found that the implemented system helped in getting a better understanding of Gabor filter responses. We think that use of a domain dependent interaction tool and additional summarization graphics may be useful in a more general Information Visualization setting.

General Terms

Gabor Filters, Information Visualization, High Dimensional Data

Keywords

Gabor Filters, Information Visualization, High Dimensional Data

1. INTRODUCTION

Spatial frequencies and their orientations are important characteristics of textures in images. Figure 2 shows examples of spatial textures with characteristic frequency and orientations. The frequency characteristics of images can be analyzed using spectral decomposition methods like Fourier analysis. We will illustrate spectral analysis for the simpler case of 1D signals. Consider the sinusoid shown in Figure 1(a). The magnitude of its Fourier spectrum is shown in Figure 1(b) - the peak corresponds to the frequency of the sinusoid. Figure 1(c) shows another sinusoid whose frequency is double that of the previous one; Figure 1(d) shows the magnitude of its spectrum. Suppose we add these two sinusoids then we will obtain a signal as shown in Figure 1(e). Doing a spectral analysis on this would show the composition of the signal - the two peaks in Figure 1(f) correspond to the component sinusoids. Fourier analysis has proven to be one of the most powerful tools in signal processing. However, a key problem with Fourier analysis is that spectral features from different parts of the image are mixed together. Many image analysis applications, e.g. object recognition, tracking, etc., require spatially localized features. Gabor filters are a popular tool for this task of extracting spatially localized spectral features [1, 2].

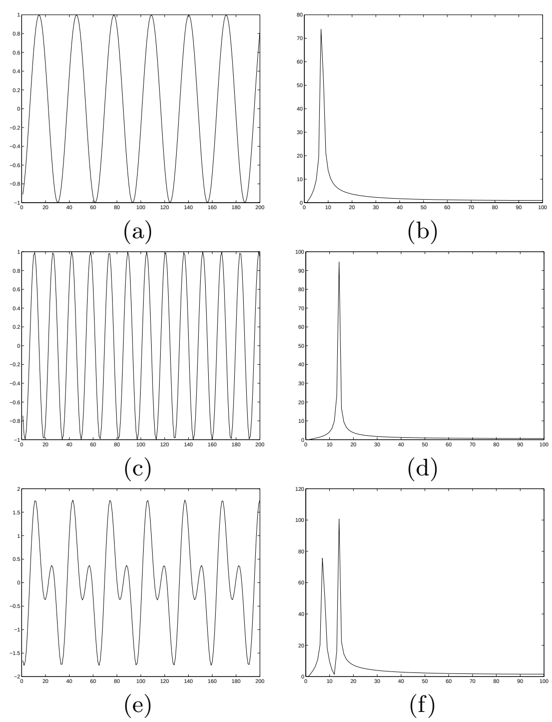

Figure 1: (a) & (b) A sinusoid and its spectrum. (c) A sinusoid with twice the frequency, (d) its spectrum. (e) Combination of the two sinusoids and (f) its spectrum.

A Gabor filter bank's response to an image consists of 4 dimensions - two of which directly correspond to the image

*The authors wish to thank Ben Shneiderman and Mustafa Bilgic for review and discussion about the work.

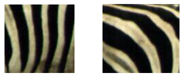

(a) (b)

Figure 2: Example of spatial frequencies in images: (a) Vertical stripes - the frequencies would have horizontal orientation, and (b) Curved stripes

plane. Visualizing a 4D space on a screen is difficult. The focus of our paper is to provide a good interactive interface for this. To give a better description of the problem, we first introduce Gabor Filters in more depth. Then we will discuss our interface and its relation to current work in Information Visualization.

1.1 Introduction to Gabor Filters

A Gabor filter is obtained by modulating a sinusoid with a Gaussian. For the case of one dimensional (1D) signals, a 1D sinusoid is modulated with a Gaussian. This filter will therefore respond to some frequency, but only in a localized part of the signal. This is illustrated in Figure 3. For 2D signals such as images, consider the sinusoid shown in Figure 4(a). By combining this with a Gaussian (Figure 4(b)), we obtain a Gabor filter - Figure 4(c). Let $g(x, y, \theta, \phi)$ be the function defining a Gabor filter centered at the origin with θ as the spatial frequency and ϕ as the orientation. We can view Gabor filters as:

$$g(x, y, \theta, \phi) = \exp(-\frac{x^2 + y^2}{\sigma^2}) \exp(2\pi\theta i(x\cos\phi + y\sin\phi))) \qquad (1)$$

It has been shown that σ, the standard deviation of the Gaussian kernel depends upon the spatial frequency to measured, i.e. θ. In our case, $\sigma = 0.65\theta$. Figure 5 shows 3D plots of some Gabor filters and the intensity plots of their amplitudes in the image plane. See [3] for an interactive tool to explore 2D Gabor filters.

The response of a Gabor filter to an image is obtained by a 2D convolution operation. Let $I(x, y)$ denote the image and $G(x, y, \theta, \phi)$ denote the response of a Gabor filter with frequency θ and orientation ϕ to an image at point (x, y) on the image plane. $G(.)$ is obtained as

$$G(x, y, \theta, \phi) = \int\int I(p, q)g(x - p, y - q, \theta, \phi)\; dp\; dq \qquad (2)$$

Consider the image of a zebra shown in Figure 6(a). If we apply a Gabor filter oriented horizontally on this image then it will give high responses wherever there are horizontal stripes present on the zebra. Figure 6(b) shows the amplitude of the response of such a horizontally oriented Gabor filter for the image.

1.2 Previous Work

The GRID principles [4] provide a general strategy for dealing with multi-dimensional data. We have used these principles here to guide our interface design. These principles would dictate that we begin to visualize our 4-dimensional space by looking at the 2-dimensional projections, and this

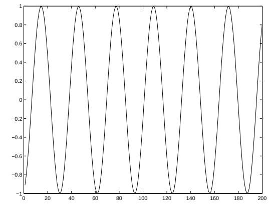

(a)

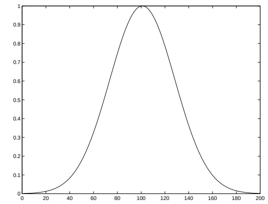

(b)

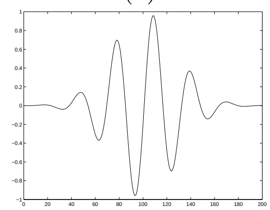

(c)

Figure 3: Gabor filter composition for 1D signals: (a) sinusoid, (b) a Gaussian kernel, (c) the corresponding Gabor filter.

has proven useful to us. The next dictate is to rank which projections are worth considering. Here, we do not need to do this dynamically, since we are always projecting the same 4 dimensions. We can therefore predict in advance which 2-dimensional projections are informative. As we will discuss in detail later, these projections are the (x, y) plane, and the (θ, ϕ) plane. We found that if the user is given these projections, the other possible projections add little. Gross et.al. present an approach for generating static visualization of Gabor filter responses using projections [5]. However, simply showing these 2-dimensional projections statically does not give a satisfactory impression of the 4-dimensional data, since many 4-dimensional spaces correspond to the same projections. We therefore included two techniques in our visualization to give a richer impression of the data. First, we designed a simple interface which allows the user to interact with the projections: the user can restrict *what parts* of the projected dimensions are visible. Second, we include additional visualizations to give information about *where* in the projected dimensions the data came from.

2. OUR APPROACH

2.1 One Dimensional Visualization

We have devised a simple way to view the responses of Gabor filters in one dimension. These filter responses can be nicely summarized in a static one-dimensional graph. This is interesting in its own right, and also provides an introduction to our approach for two dimensional filters.

Take the response of a one-dimensional Gabor filter bank to to be $G(x, \theta)$, where x is 'position' and θ is frequency. By creating an array indexed by x and θ and encoding the

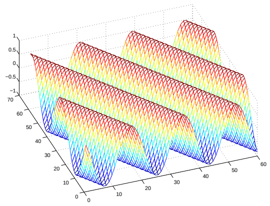

(a)

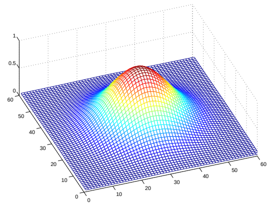

(b)

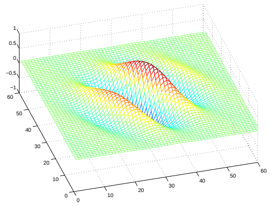

(c)

Figure 4: Gabor filter composition: (a) 2D sinusoid oriented at 30° with the x-axis, (b) a Gaussian kernel, (c) the corresponding Gabor filter. Notice how the sinusoid becomes spatially localized.

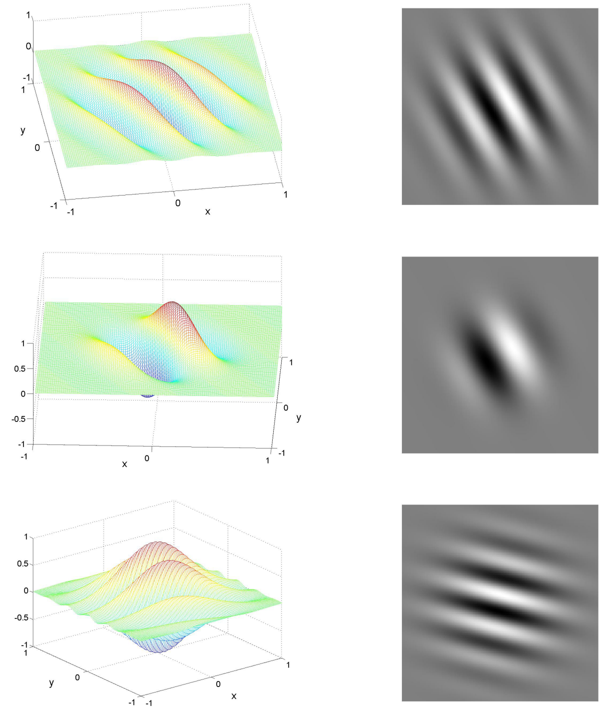

Figure 5: Example of Gabor filters with different frequencies and orientations. First column shows their 3D plots and the second one, the intensity plots of their amplitude along the image plane.

strength of the response as color, we can visualize the entire filter bank response in a single figure. For examples on synthetic signals, see Figures 7 and 8. For an example from a real signal, see see Figure 9. Observe that for real signals, it is very difficult to predict how a filter bank will respond to a given signal. This is the major motivation for this work.

2.2 Other Possibilities

A straight forward extension of the 1D visualization approach would be to simply show a matrix of intensity plots on the screen. Each intensity plot would show the amplitude of Gabor filters for a particular orientation and frequency. The frequency could vary along the row in the matrix and orientation could vary along the columns. The problem with this approach is the lack of ability of the user to interact with the filter bank parameters. Typically users like to be able to choose ranges of orientations and frequencies of the Gabor filter bank and to observe the responses over the whole image. Natural images rarely respond to specific frequencies or orientation but rather exhibit a spread over these parameters. Ability to dynamically choose the range of parameters helps in better understanding of the response characteristics. Another issue is the pragmatics of screen real estate. Typical images of interest in computer vision research are of size 300×200. Assuming that the video screen resolution is 1024×768 and we can occupy the whole screen with the intensity plots, we can show only 3 orientations and 3 scales. Even downsizing the image by half will only increase these to 7 and 7 respectively. Downsizing further might pose problems in visualization.

Medical imaging also involves visualizing images with multiple modalities simultaneously [6, 7, 8]. However, here the emphasis is on capturing the 3D human body structure. The usual approach is to stack the different image planes on top of one other and allowing the user to slice the across these planes. Notice that in our case we are dealing with 4 dimensions where only two have any explicit spatial meaning. The other dimension would be created artificially by stacking the image planes. Choosing a range of parameters would involve rotating the stack of image planes around and choosing a volume. It has been cited in visualization literature that 3D rotations during visualizations are often disorienting as it is difficult to keep track of a frame of reference over the course of interaction. In our work, we have tried to get the best possible interaction while confining ourselves to 2D visualization.

Another option would be to reuse our approach of visualizing Gabor filter responses to 1D signals. The user could be given an interface to enable him/her to slice an image into a strip. Then we could apply Gabor filters of different frequencies along this strip and stack them as shown in Figure 7. However, it would be difficult to simultaneously view responses for multiple orientations. Moreover, images have an inherently 2D structure - applying filters along 1D strips will ignore this.

Parallel coordinates are a popular approach for visualizing

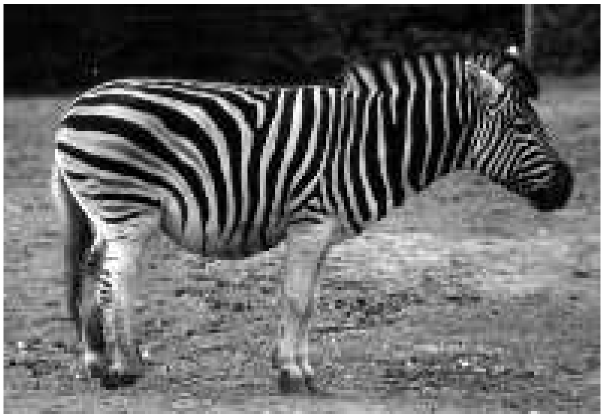

(a)

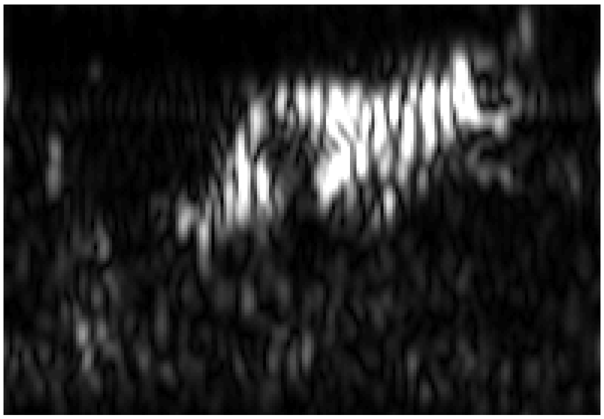

(b)

Figure 6: (a) An image, (b) The response for Gabor filter oriented horizontally - white indicates high amplitude of response, black indicates low response. Notice how regions of vertical stripes are highlighted.

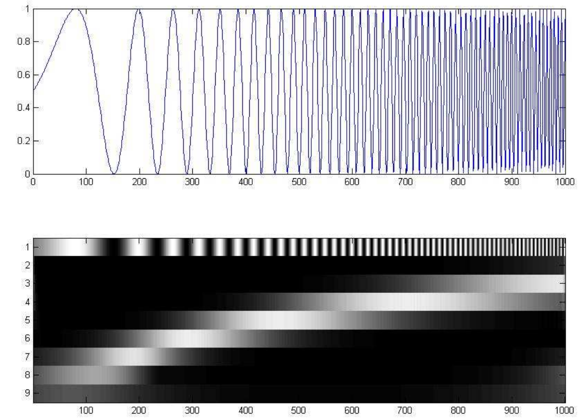

Figure 7: 1 dimensional Filter response for a synthetic signal

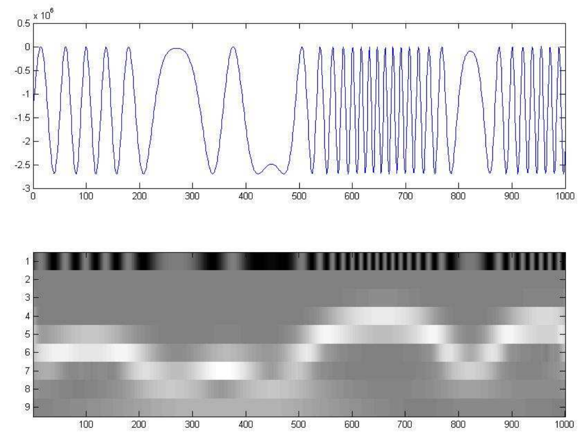

Figure 8: 1 dimensional Filter response for a synthetic signal

multi-dimensional data [9, 10]. Each dimension is plotted along an axis and all axes are placed parallel to one another. Each data point in the high dimensional space is represented by correspondingly joining the axes with line segments. In our case, two of the dimensions are coordinates on the image plane and hence plotting them on parallel coordinates might not be a good idea. (Figure 10) Star Coordinates is another popular visualization tool for multi-dimensional data [11]. However, it might not useful in our case as two of the dimensions have explicit spatial meaning.

2.3 Our Approach

Now, for a two dimensional Gabor filter bank, the situation is much more difficult. Take the response for such a bank to be $G(x, y, \theta, \phi)$, where (x, y) is the position of the filter relative to the input signal, θ is the frequency of the filter, and ϕ is the orientation of the filter. It is clear that no static presentation will allow us to view the response of the entire filter bank. Here we adopt the philosophy that, in order to give a users an understanding of the response of this 4 dimensional filter bank, *interaction* is necessary.

We feel that 2 dimensional projections are again the best way to view this data. Thus, our approach is essentially in line with the GRID principles, but applied in an unusual way. The obvious thing here, would be to observe that with 4 dimensions, there are 6 possible projections- why not simply show them and be done with it? There are two reasons.

The first and minor reason is that most of these 6 projections are not meaningful. For example, a projection onto the (y, ϕ) plane is difficult to intepret. There are really only two projections with natural interpretations: onto the (x, y) plane, and onto the (θ, ϕ) plane.

The second and more important reason, is that we have *too much data* for full projections. If we simply project the data downwards onto the (x, y) plane, we will be able to see the maximum filter response for each image point, but we will have no idea what part of the filter gave this response. We combat this problem in two ways. First, we allow users to project onto this plane, but we also allow them to *restrict*

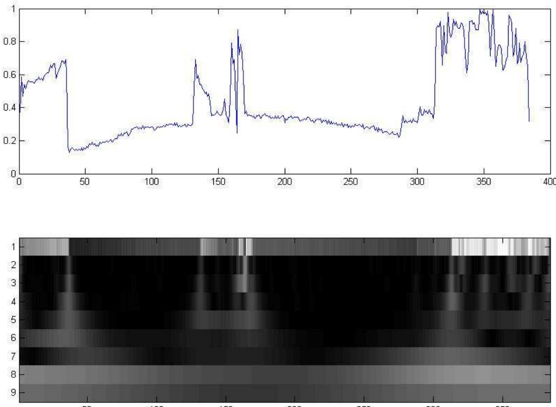

Figure 9: 1 dimensional Filter response for a real signal

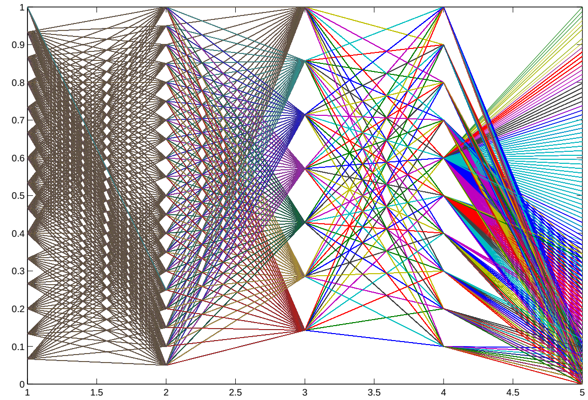

Figure 10: Parallel coordinate plot for a small 10x10 image. The dimensions, from left to right are: $x, y, \theta, \phi, G(x, y, \theta, \phi$ (Since we need to show not just the 4 dimensions, but the strength of each point therein, in some sense we are visualizing a 5-dimensional space.) Notice that even for this tiny image, Parallel Coordinates are very difficult to understand.

what portions of the other dimensions are projected. Secondly, we use additional plots to show *where* in the extra dimensions the maximum value came from.

Our interface has five plots:

1. Original image.
2. G plot: the (x, y) projection of the Gabor filter responses.
3. θ plot: this shows frequency of maximal Gabor filter response for different points on the image plane.
4. ϕ plot: this shows the orientation of maximal Gabor filter response for different points on the image plane.
5. (θ, ϕ) projection of the Gabor filter responses.

In addition, we have two interaction widgets:

1. (θ, π) interaction widget: this is shown on the (θ, ϕ) projection plane and is used to restrict the range of parameters of the Gabor filter bank. The user selects a range in the frequency and orientation dimensions of the filter bank: $(\theta_{min}, \phi_{min}) \rightarrow (\theta_{max}, \phi_{max})$. The program then finds, for every pair (x, y), the θ and ϕ such that $G(x, y, \theta, \phi)$ *is maximum*, subject to $\theta_{min} \leq \theta \leq \theta_{max}$, $\phi_{min} \leq \phi \leq \phi_{max}$. The three figures, namely (x, y) projection, θ plot and ϕ plot show G, θ and ϕ for each point on the image.
2. (x, y) interaction widget: this is shown on the original image's (x, y) plane. It is used to select a rectangular area on the image plane: $(x_{min}, y_{min}) \rightarrow (x_{max}, y_{max})$ Then, for each pair (θ, ϕ), the program finds the maximum $G(x, y, \theta, \phi)$ such that $x_{min} \leq x \leq x_{max}$, and $y_{min} \leq y \leq y_{max}$. The (θ, ϕ) projection summarizes the $G(x, y, \theta, \phi)$ for each pair (θ, ϕ).

Though the technique was created specifically for the problem of Gabor Filter visualization, we feel that our fundamental idea may be of interest to the general Information Visualization community, if stated more generally. When doing a projection of high dimensional data to a two-dimensional space, is it necessary to give up all information about the projected dimensions? We argue that it is *not* necessary and provide two techniques. First, the user can interactively restrict what portions of the larger high-dimensional space are projected. Secondly, additional plots can be used to show the values of each projected dimensions. These techniques might prove useful if applied to more general information visualization tools, such as HCE.

3. INTERFACE

3.1 Original Image

This is simply a plot of the image that the Gabor filter bank is responding to. This is important to show, because users need to compare positions on this image to positions on the other plots. The zebra image is shown in Figure 6(a).

3.2 (x, y) Interaction Widget

The user moves a box on the image to select $(x_{min}, y_{min}) \rightarrow (x_{max}, y_{max})$. This is shown in Figure 11. Notice that in this example, the user has selected a portion of the Zebra with stripes slightly off the horizontal. The program then finds the Gabor filter parameters responding to the image signal within this area. These are shown in the (θ, ϕ) projection.

3.3 (θ, ϕ) Projection

The (θ, ϕ) projection is special because the orientation, ϕ, is cyclic in the range $[0, \pi]$. In order to show this clearly, we plot the projection in the form of a disc. Here, the frequency, i.e. θ, varies along the radius, while ϕ varies along the angular direction. The plot widget is shown in Figure 12.

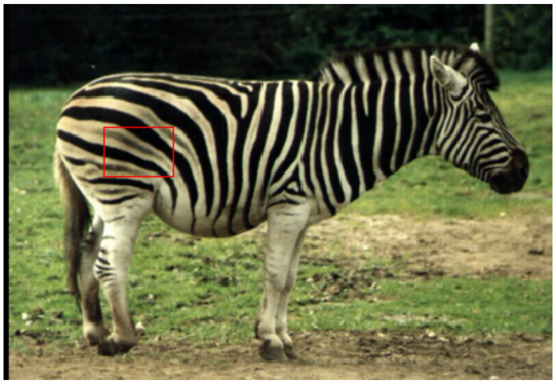

Figure 11: (x, y) Interface Widget

The disc is discretized along the radius and angular direction - in our case there were 7 quanta along θ and 10 quanta along the ϕ direction. Once the user has selected a range $(x_{min}, y_{min}) \rightarrow (x_{max}, y_{max})$, the program finds, for each pair (θ, ϕ), the maximum response, $G(x, y, \theta, \phi)$, such that $x_{min} \leq x \leq x_{max}$, and $y_{min} \leq y \leq y_{max}$. The maximum response is then plotted onto the (θ, ϕ) Projection - higher the response, brighter the intensity of the plot. Observe that the strongest filter responses, in this case, are medium-frequency, and slightly off the horizontal- this corresponds naturally to the zebra. We found that in this projection, any windows attempting to summarize the range of x and y which yielded the maximum response were not helpful, as user is generally examining a small range.

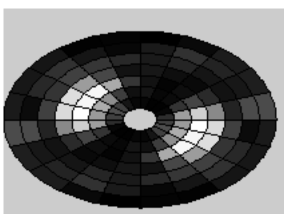

Figure 12: (θ, ϕ) Projection

3.4 (θ, ϕ) Interface Widget

This interface is based upon a simple widget which allows the users to select the range $(\theta_{min}, \phi_{min}) \rightarrow (\theta_{max}, \phi_{max})$. The widget is plotted on the (θ, ϕ) projection. It is in shape of wedge in the (θ, ϕ) projection disc. By varying its radial width, the user can choose different ranges of θ. Varying the angular width changes the range of ϕ in the Gabor filter. This scheme was chosen because it has a natural interpretation: the angles spanned by the selected region, correspond to the orientation of the Gabor filters selected, while the radius corresponds to the frequency. This is shown in Figure 13. There, the user has selected orientations close to vertical, and medium frequencies.

3.5 G plot

The plot of the $G(x, y, \theta, \phi)$ itself is shown in Figure 14. Strictly speaking this is the projection of the 4-dimensional space down into the (x, y) plane, subject to the restrictions set by $(\theta_{min}, \phi_{min}) \rightarrow (\theta_{max}, \phi_{max})$. Notice that here we can see what parts of the image have a strong response in the range set above, but this figure alone tells us little about which orientations and frequencies gave these strong responses.

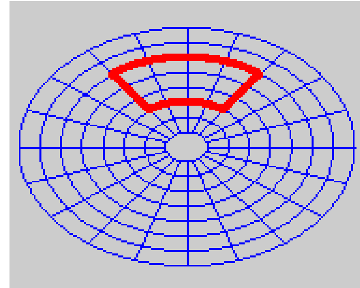

Figure 13: Interface Widget

Figure 14: G plot

3.6 ϕ plot

The plot of the ϕ found for each (x, y) such that $G(x, y, \theta, \phi)$ is maximum is shown in Figure Figure 15. This is not a projection per-se, but rather tells us about what parts of the projected dimension ϕ yielded the strongest response. Rather than coding orientation as an intensity or color, we have chosen to plot small lines with the same orientation as ϕ. This is much easier to interpret, but we can only display ϕ for every few (x, y). (Otherwise the lines cover the entire image, and one cannot see anything at all.)

3.7 θ plot

The plot of the θ found for each (x, y) such that $G(x, y, \theta, \phi)$ is maximum is shown in Figure 16. Again, this is not exactly a projection, but gives us information about how the projection was formed. Here, frequency is encoded as an intensity-darker colors correspond to higher frequencies. This encoding is not completely natural, but we were unable to find a more intuitive way to display this.

4. USING THE INTERFACE

The interface widgets can be used for exploring the filter responses. By tightly coupling the projections and auxillary plots the user's understanding of the multi-dimensional

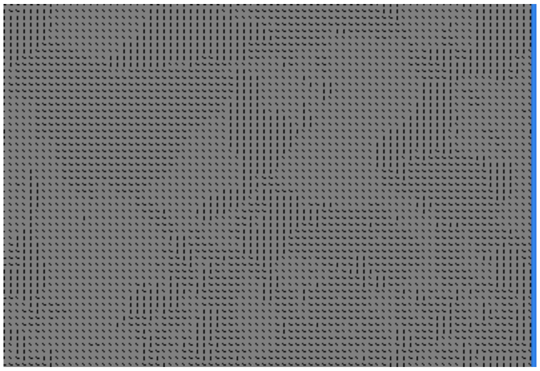

Figure 15: ϕ plot

Figure 16: θ plot

space is enhanced. The user can use the (x, y) interface widget to find out the filter responses in regions of interest in the image. He or she can then use the (θ, ϕ) interface widget to find other regions in the image which produce similar responses from the Gabor filter bank.

5. CONCLUSION

We have presented a system for the visualization of a difficult multi-dimensional space- the space of a Gabor Filter bank's response to an image. The chief difficulty in doing this is that no single, static picture can show all information about a multi-dimensional space such as this. We applied two techniques to give additional information about the projections- an interactive widget to change the visible range of the projected dimensions, and additional graphics which summarize the responses in the projected dimensions. Thus, though we view this four-dimensional space through 2-dimensional projections, we allow the user to understand all dimensions, not just the plane of projection. These techniques may be useful in a more general Information Visualization setting- when projecting to a lower-dimensional space, it is possible to retain information about all dimensions.

6. REFERENCES

[1] D. J. Gabor, "Theory of communication," *IEE*, vol. 93, no. 26, pp. 429–457, 1946.

[2] A. C. Bovik, M. Clark, and W. S. Geisler, "Multichannel texture analysis using localized spatial filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 1, pp. 57–73, 1990.

[3] "An interactive tool for visualizing Gabor filters." [Online]. Available: http://www.cs.rug.nl/ imaging/simplecell.html

[4] S. Jinwook and B. Shneiderman, "A rank-by-feature framework for interactive exploration of multidimensional data," Univ. of Maryland - College Park, Tech. Rep. HCIL-2004-31, 2004. [Online]. Available: ftp://ftp.cs.umd.edu/pub/hcil/Reports-Abstracts-Bibliography/2004-31html/2004-31.pdf

[5] M. H. Gross and R. Koch, "Visualization of multidimensional shape and texture features in laser range data using complex-valued Gabor wavelets," *IEEE Trans. Visualization and Computer Graphics*, vol. 1, no. 1, 1995.

[6] C. Rueden, K. W. Eliceiri, and J. G. White, "VisBio: A computational tool for visualization of multidimensional biological image data," *Traffic*, vol. 5, pp. 411–417, 2004.

[7] "Analyze: A software from mayo clinic for visualizing medical data." [Online]. Available: http://www.mayo.edu/bir/HTMLs/robb_FNTF_94.htm

[8] "3DViewNix: A software from the Univ. of Penn. for visualizing multimodal medical images." [Online]. Available: http://mipgsun.mipg.upenn.edu/ Vnews/info/features.html

[9] A. Inselberg and B. Dimsdale, "Parallel Coordinates: A tool for visualizing multidimensional geometry," in *IEEE Symposium on Information Visualization*, 1990, pp. 361–375.

[10] D. A. Keim, "Information visualization and visual data mining," *IEEE Trans. Visualization and Computer Graphics*, vol. 8, no. 1, pp. 1–8, 2002.

[11] E. Kandogan, "Visualizing multi-dimensional clusters, trends, and outliers using star coordinates," in *KDD '01: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 2001, pp. 107–116.

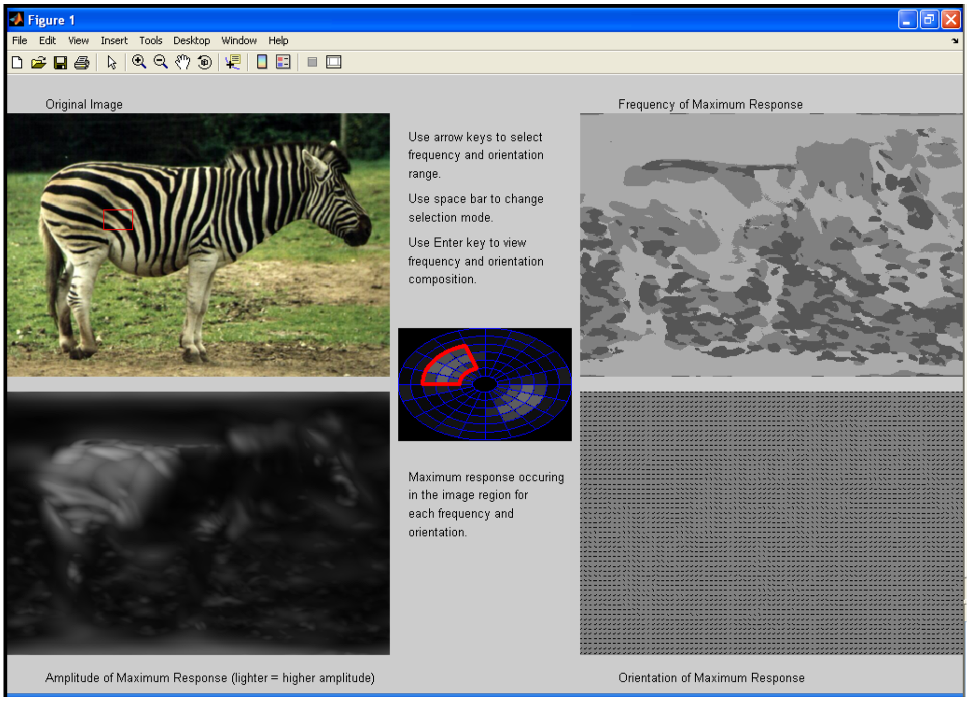

Figure 17: Layout of the interface: top left corner - original image, bottom left corner - (x, y) projection of G plot, top right corner - θ plot, bottom right corner - ϕ plot, and center - (θ, ϕ) projection of G plot. The (x, y) and (θ, ϕ) interface widgets are shown in red on the top left and center plots respectively.

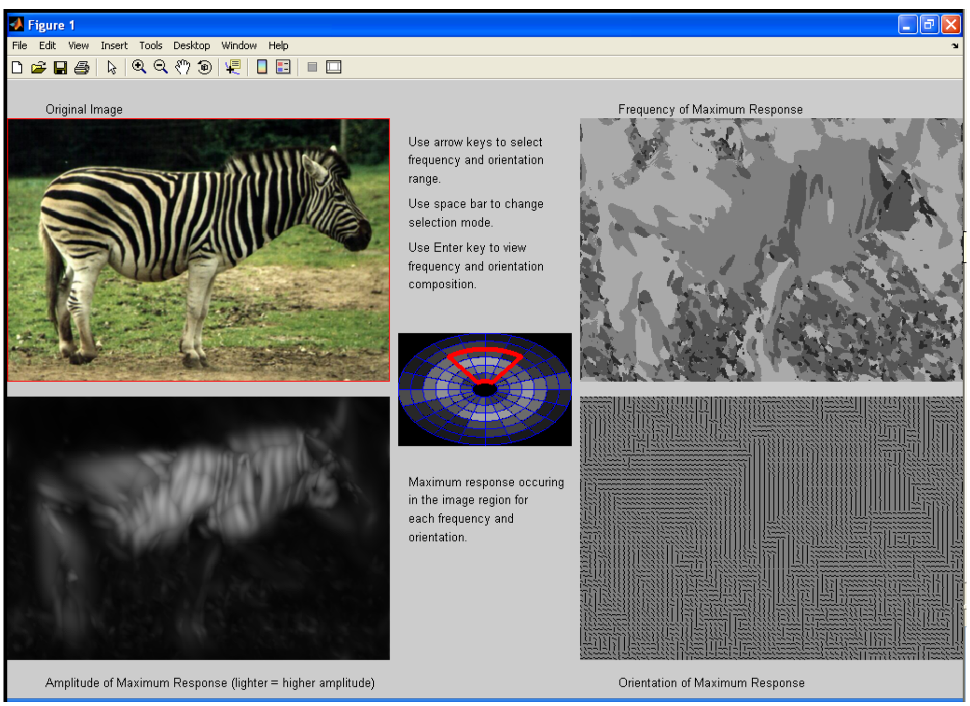

Figure 18: The user has selected the entire (x, y) range, and Gabor filters with nearly vertical orientations ϕ, with medium or high frequency θ. Notice that in the bottom left plot, parts of the zerba with vertical edges have a high response.

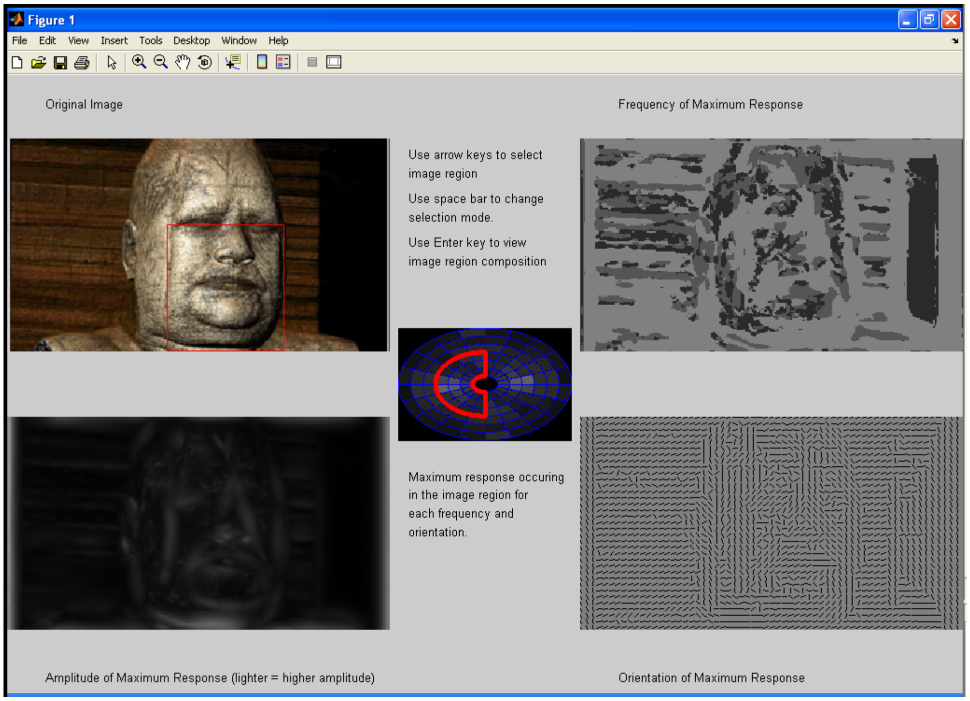

Figure 19: A different image. The user has selected the cadaver's face area. The horizontal lines in the area result in higher responses for Gabor filters with horizontal orientation, as show in the (θ, ϕ) plot in the center. Because the user has selected high frequencies with any orientation, the edges in the image are highlighted.

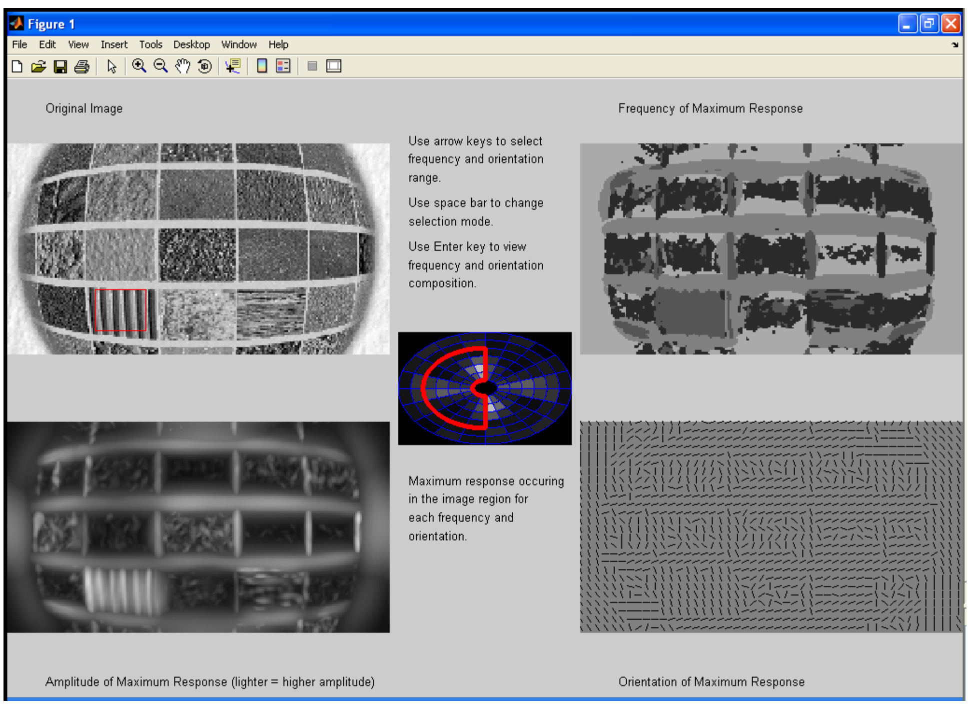

Figure 20: The user has again selected all orientations ϕ, but now with medium to high frequencies θ