

Statistics of Infrared Images

Nigel J. W. Morris
University of Toronto
nmorris@cs.toronto.edu

Shai Avidan
MERL
avidan@merl.com

Wojciech Matusik
MERL
matusik@merl.com

Hanspeter Pfister
MERL
pfister@merl.com

Abstract

The proliferation of low-cost infrared cameras gives us a new angle for attacking many unsolved vision problems by leveraging a larger range of the electromagnetic spectrum. A first step to utilizing these images is to explore the statistics of infrared images and compare them to the corresponding statistics in the visible spectrum. In this paper, we analyze the power spectra as well as the marginal and joint wavelet coefficient distributions of datasets of indoor and outdoor images. We note that infrared images have noticeably less texture indoors where temperatures are more homogenous. The joint wavelet statistics also show strong correlation between object boundaries in IR and visible images, leading to high potential for vision applications using a combined statistical model.

1. Introduction

The statistics of natural images have been extensively studied in recent years. The key observation being that images are not a random collection of pixel values and learning the statistics of the data can lead to significant advances in vision applications, as well as further our understanding of biological vision systems.

Olshausen and Field [10] suggest a sparse representation as an efficient coding of natural images and link it to the design of the biological vision system. Ruderman [14] discusses the approximate scale invariance property of natural images. Simoncelli [1] found significant dependencies of the wavelet coefficients in natural image statistics. Huang and Mumford [7] conducted a large scale experiment on the statistics of natural images that allowed them to fit mathematical models to some of these statistics and explain others in terms of local image features.

This analysis has led to significant advances in the fields of image processing and image analysis. In the field of image processing, one can consider applications such as coding, restoration, texture synthesis, de-blurring, de-noising or super-resolution [12, 16, 11], while in the field of image analysis one can consider applications such as provid-

ing context for object detection [17], or image forensics [8].

There has been some work on the statistics of other image representations and modalities. For example, Dong and Atick [2] have considered the statistics of natural time-varying images, showing that the power spectra relies non-separably both on spatial and temporal frequencies. Weiss [18] used a prior that assumes that illumination gives rise to a sparse filter output and this in turn allows him to recover both the reflectance image and the illumination images from a sequence of natural visible images. Dror *et al.* [3] studied the statistics of real-world illumination and found that they share many properties with the statistics of natural images. However, their power spectra differ significantly. This difference is attributed to strong localized light sources such as sun. Huang *et al.* [6] studied the statistics of range images. To the best of our knowledge a study of the statistics of infrared (IR) images has not been published to date.

In the next section we examine the image formation process for infrared images and explain our image acquisition process. Then we analyze the power spectra of infrared images compared to visible spectra images. Next, we examine the marginal and joint distributions of IR datasets and contrast them to the visible spectrum statistics. We also examine the joint distributions between infrared wavelet coefficients and the corresponding visible coefficients. Then we show how the marginal statistics can be re-mapped, altering the structure and texture of images. Finally we present our conclusions.

2. Infrared Imaging

Before we analyze the statistics of infrared images, it is important to understand the image formation process. In this work, when we refer to infrared, we are specifically referring to wavelengths on the electromagnetic spectrum from 4 to 12 microns in contrast to wavelengths in the visible spectrum, which roughly lie between 0.4 to 0.7 microns. This infrared wavelength spectrum is often referred to as far infrared (FIR) and we must be careful not to confuse it with near infrared (NIR) which covers wavelengths between 0.7 to 1 micron and can be captured by most CCD arrays. NIR exhibits significantly different properties from FIR and

is much closer to the visible spectrum in appearance. Although some literature further divides FIR into sub-spectra [15], we make no such distinctions since our imaging device and common thermal cameras continuously capture the entire FIR spectrum.

Infrared radiosity depends on the imaged object’s temperature, the object’s thermal emissivity as well as the reflected thermal radiance from other objects or thermal sources [9]. Several observations can be made about the appearance of infrared images. First, thermal conduction is prominently visible, as objects with different temperatures that touch each other visibly conduct and spread out energy at the point of contact. The relatively slow temporal nature of this phenomenon is also apparent, e.g. hot fingerprints remain for several seconds after the finger has left. We also note that man-made objects generally have consistent emissivity. If they are not thermal sources, their temperature becomes homogenous through conduction across the object. Second, the difference in wavelength between the visible and infrared spectra causes dramatic changes in transmission and reflection as electromagnetic waves in the infrared spectrum interact with media. Infrared waves are much more easily blocked than their visible counterparts with all but the thinnest glass sheet being impervious. In our experiments, even clear plastic CD cases allowed no infrared transmission, although transparent bubble wrap allowed most infrared energy to pass through. Due to their relatively long wavelength, surface roughness is proportionally less significant a detractor to infrared reflection. This becomes obvious as matte, tarmac pavement can act as a specular reflector of infrared radiation.

In this investigation we captured data using a single axis, multiple parameter (SAMP) camera that combined an infrared camera (Thermoteknix Miricle 110KS) with a visible light camera (Basler A601fc). Figure 1 shows how the two cameras were aligned perpendicularly with a gold dichroic beam-splitter reflecting the IR radiation from the scene towards the IR camera, while allowing the visible radiation to be transmitted to the visible light camera. The images were then further aligned using a manually specified homography. The following sections compare the statistics of sets of these image pairs. Since the cameras had different capture resolutions we scaled all images down to 320×240 pixels.

While much of the previous work on image statistics has focussed on natural scenes, infrared images tend to be more meaningful in urban environments where there is a greater thermal variation. Using the above acquisition system, we captured over 400 image pairs of indoor and outdoor scenes. We captured 70 indoor image pairs, 290 outdoor urban daytime image pairs as well as about 30 country side daytime image pairs and about 30 urban nighttime image pairs¹.

¹We will make the IR and visible image collection available online at <http://www.dgp.toronto.edu/nmorris/IR/>




Figure 1. SAMP Camera apparatus

In the following sections we focus on the indoor images and the outdoor urban daytime images.



Figure 2. Exemplar image pairs taken from our dataset

3. Power Spectra Analysis

In the 50’s, television engineers discovered the power law for image power spectrum decay. It says that the spectral power decays at the rate of: $1/f^{2-\eta}$ where f is the spa-

tial frequency and η depends on the image but is generally small. This analysis method has been applied to natural images [4, 13] and found to fit both the power spectra of both the logarithm of image intensities as well as the raw image intensities. We analyzed the average spatial power spectra for each of our four datasets using both the raw intensities and the logarithm of the intensities.

For the raw intensities, we found that, while the power law fits most of our data well, the DC component does not fit well. This effect is especially noticeable in the IR spectrum and can be explained by the relative lack of texturing as we would expect in infrared images. Since our visible images are strictly of urban scenes rather than natural environments, that would also explain an increased DC component in the visible spectrum. We found that a generalized Laplace distribution gave a better fit for the whole power spectrum where the power can be described as:

$$\mathcal{P}(x) \propto \exp(-|x/s|^\alpha). \quad (1)$$

Figure 3 shows the raw intensity spatial power spectra on log-log axes along with the Laplace distribution fits for each dataset. There is a small local maximum close to the highest frequency in the IR power spectrum. This is caused by structured noise in our infrared imaging device and appears in almost every image.




Figure 3. Spatial power spectra for raw image intensities. The dotted lines show the fitted Laplacian distributions.

We found that the power spectrum for the log of image intensities of the visible dataset fit the power law well, but the IR data fit was modeled more effectively by a Gaussian. There are noticeable non-linearities at the highest and lowest spatial frequencies that do not fit the power law (Figure 4).




Figure 4. Spatial power spectra for log image intensities. The dotted lines show the fitted Gaussian distribution for the IR power spectrum on the left and the fitted power law for the visible spectrum on the right.

4. Wavelet Statistics

Wavelet coefficient statistics are a powerful tool for measuring the characteristics of images. Distributions of coefficients at various scales and orientations provide us with statistical dependencies that can be leveraged in many applications [3, 12, 16]. In order to better approximate the human visual system as in previous work [3, 6], we use the log of the image intensities in our wavelet analysis. We use ‘Haar’ wavelets as commonly used in the natural image statistics literature [3, 6].

All figures of wavelet coefficient distributions show the log probability since this reveals the non-Gaussian nature of the distributions. All distributions shown have been normalized before applying the logarithm so that the area under the curve sums to one.

4.1. Marginal distributions

We first present the results and analysis of the marginal wavelet distributions. Figure 5 shows a comparison of the marginal statistics of wavelet coefficients at three scales for the four image groups. We fit each distribution to a generalized Laplace distribution (Equation 1) as described in [7, 16]. We estimate the parameters using least squares error as in [7]. The estimated values for the fitted Laplace distribution parameters α and s are also shown at the top right of each histogram. The visible spectrum distributions fit very well to the Laplace curves as expected. The IR distributions also closely fit the Laplace distribution although there is some discrepancy around the mean. It has been reported [16] that natural images are generally fit with an exponent parameter α in the range of [0.5, 0.8]. The visible

spectrum distributions tend to fit this prediction with the indoor distribution exponents just under 0.5. We noted that IR distribution exponent parameters tend to fall in the range of $[0.25, 0.5]$. The IR distribution curves exhibit significantly stronger kurtosis and have more predominant peaks than the corresponding visible spectrum fits. Recall that this is consistent with our observation in Section §2 that man-made objects have homogenous emissivity and texture that would appear in the visible spectrum does not show up in the infrared images. For example, indoor scenes tend to have more homogenous temperatures since indoor air temperature is regulated and conduction reduces the temperature variation within the scene. On the other hand, outdoor images also tend to cross a larger scale, capturing objects with more diverse emissivity and temperature. This would suggest that the marginal outdoor images statistics should have less kurtosis and have heavier tails than indoor image distributions. However, weather conditions also play a major role in the statistics of outdoor scenes. The relatively high kurtosis in the outdoor marginal statistics can be attributed to cool, overcast weather conditions during capture. This weather would negate many of the textural effects caused by shadows and infrared reflectance of the sun.

4.2. Joint distributions

Another set of statistical measures commonly used to examine natural images is joint distributions of wavelet coefficients between orientations, scales and space. In Figure 6 the same set of joint distributions shown in [7, 3] are displayed alongside each other for all four sets of data (the final distribution is omitted due to space constraints). The figure shows contour plots of the log probability for each joint distribution. Horizontal, vertical and diagonal component labels indicate the three primary orientation bands from the wavelet decomposition. Each pixel in the data has a wavelet coefficient response at each orientation and the top two rows in Figure 6 show joint distributions of these responses for pairs of orientations. The histogram for horizontal orientation coefficients versus vertical orientation coefficients is computed as follows:

$$\mathcal{H}(W_l^h(i, j), W_l^v(i, j)), \quad (2)$$

where $W_l^h(i, j)$ is the wavelet coefficient value for the horizontal orientation band at scale l and pixel position (i, j) and $W_l^v(i, j)$ is the corresponding coefficient on the vertical orientation band.

We also show joint distributions of spatial relationships where neighboring pixels are indicated as ‘brother’. For example the distribution between the horizontal component and its ‘left brother’ indicate a joint spatial distribution between the set of horizontal wavelet coefficients and the set

of horizontal coefficients one grid place to the left:

$$\mathcal{H}(W_l^h(i, j), W_l^h(i - 1, j)). \quad (3)$$

The ‘parent’ relationship indicates a joint distribution with the corresponding coefficient one step up in scale within the wavelet decomposition pyramid:

$$\mathcal{H}(W_l^h(i, j), W_{l+1}^h(i, j)). \quad (4)$$

The fact that the joint statistics of the visible spectrum is not Gaussian is mirrored by the corresponding infrared statistics with all the striking polyhedral shapes also occurring in the IR distributions. Huang and Mumford [7] show that the non-Gaussian features of the joint distributions correspond to the presence of simple geometric features present in natural images. It is not surprising that the same features are also present in infrared images, since object discontinuities are generally preserved while object textures are generally not. We therefore suggest that these joint statistical structures are caused by object shapes rather than surface texture.

There is a noticeable difference in scale between the IR joint distributions and the visible ones. This is caused by the relatively small contrast between objects in the infrared spectrum. Generally, objects have small temperature differences resulting in smaller wavelet coefficient responses than the corresponding intensity changes in the visible spectrum images. If the infrared images are normalized, then the scale matches the visible images more closely. However, this introduces significant noise and we have decided to present the raw image data.

4.3. Joint infrared and visible statistics

Finally we present results of wavelet statistics of the IR coefficients directly compared to the aligned coefficients in the visible spectrum (Figure 7). We computed these statistics using a subset of 141 of the outdoor urban images. These distributions take the following form:

$$\mathcal{H}(W_l^{IR}(i, j), W_l^{Vis}(i, j)), \quad (5)$$

where $W_l^{IR}(i, j)$ is the wavelet coefficient in the IR image and $W_l^{Vis}(i, j)$ is the corresponding coefficient at the same image position in the paired visible spectrum image. The figure, as in the previous section, shows joint distributions across different orientation bands, scales and also in space. The bottom row also shows the distribution between the IR and visible coefficients for the same orientation and scale. The IR axis has been scaled to compensate for the reduced contrast and to facilitate comparison between distributions. The plots show that there is a high correlation between wavelet coefficients in the direct orientation comparisons as well as the inter-orientation, spatial and scale distributions.




Figure 5. Marginal log distributions of horizontal wavelet coefficients for a) indoor infrared images, b) indoor visible, c) outdoor urban infrared and d) outdoor visible. Three scales of coefficients from coarsest (top) to finest (bottom) are shown along with the corresponding fitted general Laplacian distributions. The estimated parameters for the distributions are shown in order at the top right of each histogram.

5. Wavelet distribution re-mapping

One of the advantages of working with the wavelet statistics of an image is that it becomes relatively easy to remap the image to a new set of distributions using a pyramid type scheme as in [5]. For instance, this allows us to map a lightly-tailed distribution to a heavily-tailed one. In order to do the re-mapping we need to perform the following steps:

1. Decompose the image into its wavelet components
2. For each component, W , at each pyramid level:
 - Create a histogram, \mathcal{H}_1 , of the wavelet coefficients
 - Create the CDF, C_1 , from \mathcal{H}_1
 - For a given target histogram, \mathcal{H}_2 , create another CDF, C_2
 - For each coefficient $W(j)$ at image position j :
 - Find $\mathcal{P}(W(j))$ from $C_1(W(j))$
 - The new coefficient value $W'(j) = C_2^{-1}(\mathcal{P}(W(j)))$
3. Reconstruct the new image using the modified wavelet components W'

Instead of a target set of histograms \mathcal{H}_2 , we can fit \mathcal{H}_1 to our distribution model (1) and then remap to an adjusted version of \mathcal{H}_1 . This is done by fitting the general Laplace distribution parameters α and s to \mathcal{H}_1 , then rescaling them to obtain a new distribution \mathcal{H}' . Then we can merely replace \mathcal{H}_2 with \mathcal{H}' and continue to step 2 of the above algorithm. Figure 8 illustrates the results of this modified version of our re-mapping algorithm on a single infrared image. It shows the re-mapped images after the fitted distribution parameters α and s have been rescaled by r_α and r_s . The α parameter adjusts the peakiness of the distribution while the s parameter provides a more global scaling of the curve. The effect of a heavy tailed distribution can clearly be seen in this figure, causing the image details to stand out.

6. Discussion and Conclusions

We have investigated the statistics of IR images and made a several observations. We found that IR power spectra, as opposed to visible image power spectra, are better fitted with a generalized Laplace curve. This is primarily caused by the lack of texture in IR images. As expected, we found that our analysis of the marginal wavelet statistics confirmed this observation. It was also evident that IR images are less useful in scenes where temperature variation is low. This is often true for indoor scenes and can be true for outdoor scenes where weather may be a factor. On the other hand, the joint statistics of the wavelet coefficients are very similar to those reported for visible images, albeit on a smaller scale, probably due to lower contrast in IR images. The close matching of joint statistics to the visible spectrum suggests that IR images do indeed capture the shape of objects while discarding much of their texture. This suggest that IR images may closely resemble intrinsic images and provide a computationally free method for obtaining immediate intrinsic images for use with other vision algorithms. The results from our investigation into the joint statistics of visible and IR images also highlight the strong correlation between object boundaries in IR and visible images.

There appears to be high potential for the application of IR image statistics, especially when combined with an aligned visible light camera. The absence of object texture in IR images should improve image de-noising algorithms and enhance compression. Supplementing the visible spectrum with the additional statistical information in the infrared spectrum should also benefit statistics-based segmentation and categorization algorithms.

7. Acknowledgements

This work was supported in part by the Province of Ontario under the OGSST program.

References

- [1] R. W. Buccigrossi and E. P. Simoncelli. Image compression via joint statistical characterization in the wavelet domain.




Figure 6. Joint log distributions of wavelet coefficients

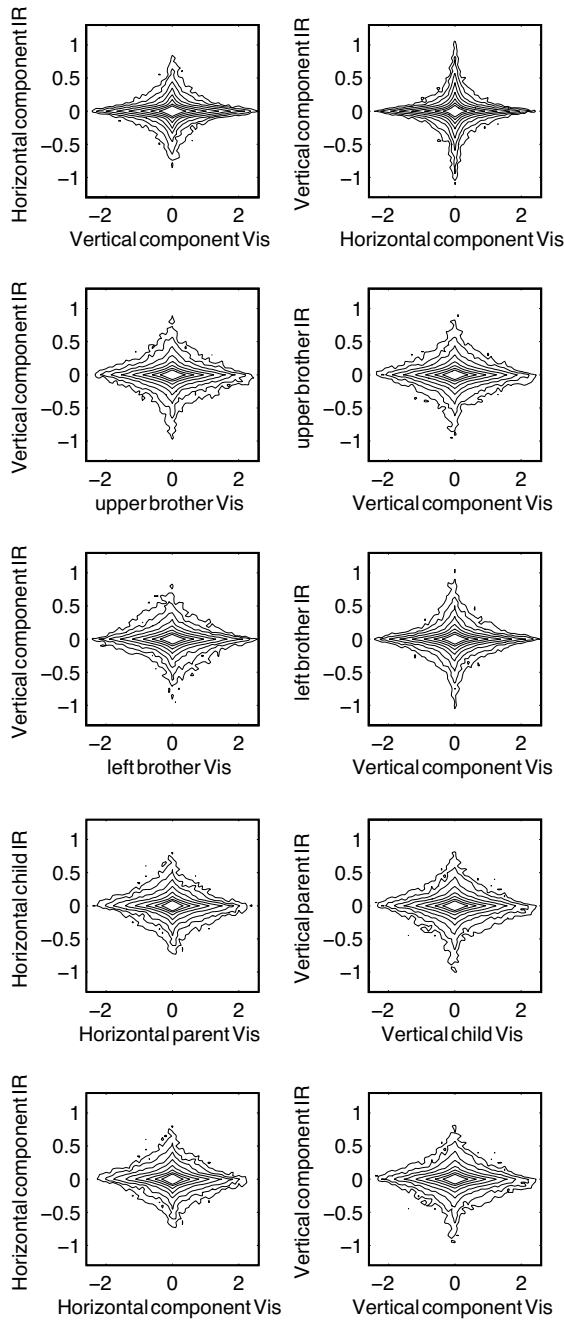


Figure 7. Joint infrared wavelet coefficient versus visible coefficient distributions.

IEEE Trans Image Proc, 8(12):1688–1701, Dec. 1999. **1**

- [2] D. W. Dong and J. J. Atick. Statistics of natural time-varying images. *Network*, 6(3):345–358, August 1995. **1**
- [3] R. O. Dror, T. K. Leung, E. H. Adelson, and A. S. Will-sky. Statistics of real-world illumination. In *CVPR (2)*, pages 164–171, 2001. **1, 3, 4**
- [4] D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *J. Optical Society of America*, 4(12):2379–2394, 1987. **3**




Figure 8. Top rows: Re-mapped images and marginal distribution curves for increasing values of r_α with $r_s = 1$. Bottom rows: Re-mapped images and marginal distribution curves for increasing values r_s with $r_\alpha = 1$.

- [5] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. In *SIGGRAPH '95*, pages 229–238, 1995. **5**
- [6] J. Huang, A. B. Lee, and D. Mumford. Statistics of range images. In *CVPR*, pages 1324–1331, 2000. **1, 3**
- [7] J. Huang and D. Mumford. Statistics of natural images and models. In *CVPR*, pages 1541–1547, 1999. **1, 3, 4**
- [8] S. Lyu and H. Farid. Detecting hidden messages using higher-order statistics and support vector machines. In *5th International Workshop on Information Hiding*, 2002. **1**
- [9] N. Nanhakumar and J. K. Aggarwal. Integrated analysis of thermal and visual images for scene interpretation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 10(4):469–481, 1988. **2**
- [10] B. Olshausen and D. Field. Natural image statistics and efficient coding. *Network*, 7:333–339, 1996. **1**
- [11] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *IJCV*, 40(1):49–70, 2000. **1**
- [12] J. Portilla and E. P. Simoncelli. Image restoration using gaussian scale mixtures in the wavelet domain. In *ICIP (2)*, pages 965–968, 2003. **1, 3**
- [13] D. Ruderman. The statistics of natural images. *Network: Comput. Neural Syst.*, 5:517–548, 1994. **3**
- [14] D. Ruderman. Origins of scaling in natural images. *Vision Research*, 37(23):3385–3395, 1997. **1**
- [15] D. Scribner, P. Warren, and J. Schuler. Extending color vision methods to bands beyond the visible. *Mach. Vision Appl.*, 11(6):306–312, 2000. **2**
- [16] A. Srivastava, A. B. Lee, E. P. Simoncelli, and S.-C. Zhu. On advances in statistical modeling of natural images. *J. Math. Imaging Vis.*, 18(1):17–33, 2003. **1, 3**
- [17] A. Torralba and A. Oliva. Statistics of natural image categories. *Network*, 14:391–412, 2003. **1**
- [18] Y. Weiss. Deriving intrinsic images from image sequences. In *ICCV*, pages 68–75, 2001. **1**