

Feature Based Patch Image Denoising Model

M Ramana Kumar

Department of Computer Sciences,
Rayalaseema University,
Kurnool, Andhra Pradesh 518002, India.

Prof Bhaskar Reddy

Department of Computer Sciences,
Sri Krishnadevaraya University,
Anantapur, Andhra Pradesh 515003, India.

ABSTRACT

Image denoising has been an active field of research with literature dating back to the 1970s. However, given the importance of the problem, considerable effort still continues to be channeled to bettering the state-of-the-art. Surprisingly, performance improvement in recent years has been somewhat limited. In this thesis, we first study the possible causes of such restricted improvement. To do so we analyze the problem of image denoising in a statistical framework. Since the best performing methods are feature based patch selection model, we frame the problem of denoising as that of estimating the underlying image patches from their noisy observations. This chapter introduces the problem of image denoising, and discusses the various sources and characteristics of noise corrupting images and why denoising is an important problem. Finally, we discuss the state-of-the-art in image denoising and its improvement based on feature based patch selection denoising model.

INTRODUCTION

1. Digital Image

A digital image is a numeric representation (normally binary) of a two-dimensional image. Depending on whether the image resolution is fixed, it may be of vector or raster type. By itself, the term "digital image" usually refers to raster images or bitmapped images [1-5].

1.1 Raster

Raster images have a finite set of digital values, called picture elements or pixels. The digital image contains a fixed number of rows and columns of pixels. Pixels are the smallest individual element in an image,

holding quantized values that represent the brightness of a given color at any specific point. Typically, the pixels are stored in computer memory as a raster image or raster map, a two-dimensional array of small integers. These values are often transmitted or stored in a compressed form [7].

Raster images can be created by a variety of input devices and techniques, such as digital cameras, scanners, coordinate-measuring machines, seismographic profiling, airborne radar, and more. They can also be synthesized from arbitrary non-image data, such as mathematical functions or three-dimensional geometric models; the latter being a major sub-area of computer graphics [9]. The field of digital image processing is the study of algorithms for their transformation.

1.1.1 Raster file formats

Most users come into contact with raster images through digital cameras, which use any of several image file formats. Some digital cameras give access to almost all the data captured by the camera, using a raw image format. The Universal Photographic Imaging Guidelines (UPDIG) suggests these formats be used when possible since raw files produce the best quality images. These file formats allow the photographer and the processing agent the greatest level of control and accuracy for output. Their use is inhibited by the prevalence of proprietary information (trade secrets) for some camera makers, but there have been initiatives such as OpenRAW to influence manufacturers to release these

Cite this article as: M Ramana Kumar & Prof Bhaskar Reddy, "Feature Based Patch Image Denoising Model", International Journal & Magazine of Engineering, Technology, Management and Research, Volume 5 Issue 2, 2018, Page 224-232.

records publicly. An alternative may be Digital Negative (DNG), a proprietary Adobe product described as “the public, archival format for digital camera raw data”.^[1] Although this format is not yet universally accepted, support for the product is growing, and increasingly professional archivists and conservationists, working for respectable organizations, variously suggest or recommend DNG for archival purposes [11].

1.2 Vector

Vector images resulted from mathematical geometry (vector). In mathematical terms, a vector consists of point that has both direction and length.

Often, both raster and vector elements will be combined in one image; for example, in the case of a billboard with text (vector) and photographs (raster).

1.3 What Is Digital Image Processing?

An image may be defined as a two-dimensional function, $f(x, y)$, where x and y are spatial (plane) coordinates, and the amplitude of at any pair of coordinates (x, y) is called the intensity or gray level of the image at that point. When x, y , and the amplitude values of f are all finite, discrete quantities, we call the image a digital image. The field of digital image processing refers to processing digital images by means of a digital computer. Note that a digital image is composed of a finite number of elements, each of which has a particular location and value. These elements are referred to as picture elements, image elements, pels, and pixels. Pixel is the term most widely used to denote the elements of a digital image. Vision is the most advanced of our senses, so it is not surprising that images play the single most important role in human perception. However, unlike humans, who are limited to the visual band of the electromagnetic (EM) spectrum [10], imaging machines cover almost the entire EM spectrum, ranging from gamma to radio waves. They can operate on images generated by sources that humans are not accustomed to associating with images. These include ultra-sound, electron microscopy, and computer-generated images. The area of image analysis (also called image understanding) is in between image

processing and computer vision. A low-level process is characterized by the fact that both its inputs and outputs are images. Mid-level processing on images involves tasks such as segmentation (partitioning an image into regions or features), description of those features to reduce them to a form suitable for computer processing, and classification (recognition) of individual features. A mid-level process is characterized by the fact that its inputs generally are images, but its outputs are attributes extracted from those images (e.g., edges, contours, and the identity of individual features). Finally, higher-level processing involves “making sense” of an ensemble of recognized features, as in image analysis, and, at the far end of the continuum, performing the cognitive functions normally associated with vision and, in addition, encompasses processes that extract attributes from images, up to and including the recognition of individual features. As a simple illustration to clarify these concepts, consider the area of automated analysis of text. The processes of acquiring an image of the area containing the text, preprocessing that image, extracting (segmenting) the individual characters, describing the characters in a form suitable for computer processing, and recognizing those individual characters are in the scope of what we call digital image processing [8].

1.4 Components of an Image Processing System




Fig.1.4. Components of a general purpose Image Processing System

As recently as the mid-1980s, numerous models of image processing systems being sold throughout the world were rather substantial peripheral devices that attached to equally substantial host computers. Late in

the 1980s and early in the 1990s, the market shifted to image processing hardware in the form of single boards designed to be compatible with industry and hard buses and to fit into engineering workstation cabinets and personal computers. In addition to lowering costs, this market shift also served as a catalyst for a significant number of new companies whose specialty is the development of software written specifically for image processing. Although large-scale image processing systems still are being sold for massive imaging applications, such as processing of satellite images, the trend continues toward miniaturizing and blending of general-purpose small computers with specialized image processing hardware. Figure shows the basic components comprising a typical general purpose system used for digital image processing. The function of each component is discussed in the following paragraphs, starting with image sensing. With reference to sensing, two elements are required to acquire digital images. The first is a physical device that is sensitive to the energy radiated by the feature we wish to image. The second, called a digitizer, is a device for converting the output of the physical sensing device into digital form. For instance, in a digital video camera, the sensors produce an electrical output proportional to light intensity. The digitizer converts these outputs to digital data [6].

Specialized image processing hardware usually consists of the digitizer just mentioned, plus hardware that performs other primitive operations, such as an arithmetic logic unit (ALU) [4], which performs arithmetic and logical operations in parallel on entire images. One example of how an ALU is used is in averaging images as quickly as they are digitized, for the purpose of noise reduction. This type of hardware sometimes is called a front-end subsystem, and its most distinguishing characteristic is speed. In other words, this unit performs functions that require fast data throughputs (e.g., digitizing and averaging video images at 30 frames) that the typical main computer cannot handle. The computer in an image processing system is a general-purpose computer and can range from a PC to a supercomputer. In dedicated applications, sometimes

specially designed computers are used to achieve a required level of performance, but our interest here is on general-purpose image processing systems. In these systems, almost any well-equipped PC-type machine is suitable for offline image processing tasks. Software for image processing consists of specialized modules that perform specific tasks. A well-designed package also includes the capability for the user to write code that, as a minimum, utilizes the specialized modules.

LITERATURE REVIEW

2.1 State Of The Art

Currently, in the most promising approaches to feature recognition and classification, a feature is assumed to consist of several parts which can be modelled more or less independently. Although this is not always likely to hold, it facilitates computation greatly without degrading the results. The various approaches differ mainly in the way the feature parts are modeled, the possible incorporation of the spatial context, the type of features used, and in the limitations imposed on the data. We will compare the results of the most successful approaches to the results of our method in the Chapter 8. Recently, [Opelt & Pinz+ 06] presented an approach where weak hypotheses are combined using boosting. As features, local descriptions of regions of discontinuity and of homogeneity are used. Regions of discontinuity are detected by several interest point detectors, while for determining regions of homogeneity the authors propose a new segmentation algorithm. The approach works with weak supervision during training, and the authors present very good results for recognition of features in complex images [2].

It is not accounted for spatial relationships in this method. Another approach to feature recognition was proposed by [Deselaers & Keysers+ 05a], where image patches are clustered using the EM algorithm [1] for Gaussian mixture densities and images are represented as histograms of the patches over the (discrete) membership to the clusters. A log-linear maximum entropy classifier is used to classify the images. The patches represent both regions of high variance, using an

interest point detector, and homogeneous regions using a regular grid projected onto the image. This method proved to perform well on a variety of databases, while also only weak supervision is necessary during training. In [Deselaers & Keysers+ 05b], the authors propose among several other improvements an extension towards fuzzy histograms to reduce discretization effects. The method presented here is partly based on this approach. In particular, image patches are used as features, too. In contrast, it does not consider spatial information, which the proposed method does. Patches are also regarded in [Paredes & Perez-Cortes+ 01], where they are classified by a nearest neighbor based voting scheme.

IMAGE PATCHES

3. Image Patches

Patches have been successfully used in several approaches to feature recognition before, for example in [Paredes & Perez-Cortes+ 01] and [Deselaers & Keysers+ 05a]. This chapter gives a definition of patches and a motivation why they are used in feature recognition. Afterwards, the process of extracting them from images is presented, which includes determining the location and size of the patches, and several preprocessing steps. As preprocessing steps dimensionality reduction, normalization with respect to brightness, and including derivatives are considered.

3.1 Patches in Feature Recognition with image patches

We denote squared subimages extracted from an image. Three parameters determine a patch uniquely: the horizontal and vertical location within the image, specified by the x- and y-coordinate of the patch center, and its size. For a given location and size, the patch can be extracted by simply determining which image pixels are located within that particular square. Patches belong to the category of local features which means that they describe properties of a certain region of an image. In contrast to that, global features provide information about an image as a whole. Typical global features used in feature recognition or in the related field of image retrieval are for example texture information, the color

distribution, or simply all pixels of the image itself. A description and evaluation of global features can be found in [Deselaers 03] and [Deselaers & Keysers+ 04]. As global features account for whole images, they tend to become inadequate if only a small portion of the image is relevant. In the domain of feature recognition, it is often the case that images have to be classified based on features which make up only a very limited part of the image. Such images are called “complex images” or “complex scenes”. As an example, consider the image in Figure 3.1. Due to the bicycle leaned against the wall, the system should answer “yes” for the question whether the image contains a bicycle or not. But obviously, the characteristic content of this image is the wall of the house with several windows, and the street in front of it.

Global features would not differ strongly if the bicycle was not present in the image at all, i.e. if the image showed the wall without the bicycle in front. In complex images, the information global features provide is not sufficient and therefore they are not well suited in this context. Local features like patches are better suited for complex images, because they represent restricted regions of the image. In this work patches are used exclusively.



Figure 3.1: A complex bicycle image

Beneficial properties of local features are:

- **Inherent translation invariance:** using patches it does not matter where in the image a certain feature is shown. An feature may be successfully detected even if it uses to be located at different locations in different images.

3.2 Extraction of Patches

In this section, the process of extracting a set of patches from a given image is studied.

3.2.1 Extraction Points

In the simplest case, a patch is extracted around each pixel of an image at a given scale. Although this seems to be a good idea, as we can be sure not to discard any parts of the image, it turns out that the resulting amount of patches is not feasible for the methods we apply. Therefore, to use a moderate number of patches per image, a subset of points needs to be chosen around which patches are extracted. In the following, different methods to determine such a subset of extraction points are presented:

- **Grid points:** choosing grid points is a trivial way of determining extraction points. A regular grid is projected onto the image which immediately gives the extraction points. With this method the extraction points are distributed uniformly over the image, which means that they are contained in homogeneous regions and in regions with high variance. The size of the gaps between the patches or whether they overlap depends on the patch size and on the distance between the extraction points which is determined by the size of the grid. In particular, the grid can be chosen such that the patches are aligned.

3.2.2 Patch Size

As stated above, a difference between the Difference-of-Gaussian points and the other extraction points is that for the former ones a patch size can be determined which is likely to perform best. For grid, random, and wavelet-based salient interest points instead the patch size has to be chosen arbitrarily. For the images of the databases presented in Chapter 7 we encountered for example a patch size of 11×11 pixels to perform well. If we assume that the features appear across all images at roughly the same size, then we can extract all patches at the same chosen patch size. Anyway, this assumption is unlikely to hold in many cases. A possibility to address this scale difference of the features is to extract the patches at different scales. In our experiments, patches are extracted at sizes of 7×7 , 11×11 , 21×21 and 31×31

pixels. The sizes are chosen to represent small, middle and large feature parts. Surely, extracting patches at multiple sizes increases the amount of data to deal with. Still, we expect the results to improve using a combination of these patch sizes, as this effect has already been studied in [Deselaers&Keysers+ 05b], where a significant improvement of the classification result using multiple patch sizes is reported.

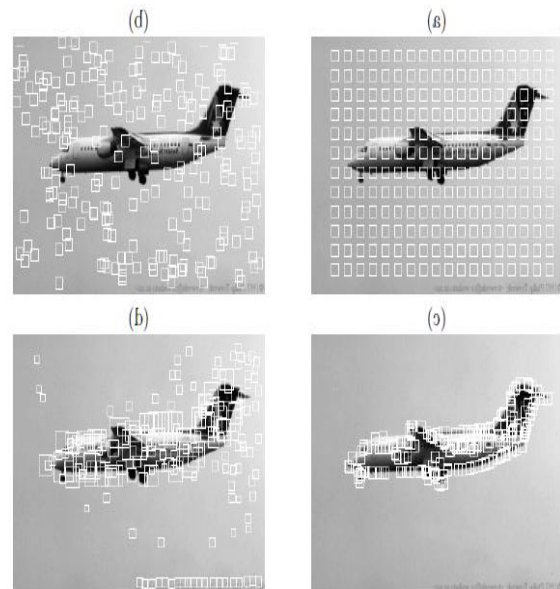


Figure 3.3: (a): Grid extraction points, (b): Random extraction points, (c): Wavelet-based salient extraction points, (d): Difference-of-Gaussian extraction points.

3.3 Patch Preprocessing

In the preprocessing phase, the extracted patches are manipulated and turned into feature vectors. Formally, for a given image X and its extracted patches, a set of feature vectors $\{x_l^L\}$ is generated, where x_l is the l -th feature vector, corresponding to the l -th extracted patch.




Figure 3.3: Untransformed patch (left), PCA transformed and backtransformed patch with 40 PCA coefficients (right).

3.3.1 Feature Reduction

In the simplest case, the pixel values of the patches can be used without any further processing as components of the feature vectors. For an $n \times n$ patch, there are n^2 gray-level pixel values and thus we obtain feature vectors with n^2 components. Even for patches of moderate size, for example 20×20 pixels, this leads to huge feature vectors. Therefore, a feature- or dimensionality reduction of the feature vectors is desirable, which is a well known problem in pattern recognition [Duda& Hart+ 01]. A commonly used reduction method is the Principle Component Analysis (PCA). PCA is an unsupervised approach to extract the appropriate features from the patches. First, the n^2 -dimensional mean vector μ and the $n^2 \times n^2$ -dimensional covariance matrix Σ are computed from all $n \times n$ patches of the training images.

$$x' = A(x - \mu)$$

where x' denotes the reduced feature vector. For most experiments described in this work, we set k to 40, as this proved to be a suitable value [Paredes & Keysers+ 02]. Figure 3.3 shows that keeping 40 coefficients is sufficient to maintain most of the information a patch contains. Note that other feature reduction techniques exist, for example the Linear Discriminant Analysis (LDA). As previous work [Kölsch 03], [Kölsch&Keysers+ 04] showed that PCA outperforms LDA in image classification using local features, in this thesis PCA is used exclusively.

3.3.2 Brightness Normalization

Another type of preprocessing that can be applied to patches is a brightness normalization. As long as not “artificial” images or images which have already been normalized beforehand are used, it is quite normal that the features appear in different images under different lighting conditions. Consider the two faces shown in Figure 3.4. While the face on the left is quite bright, the face to the right appears very dark. It seems natural to apply a brightness normalization to achieve a more homogeneous brightness of the features, which can either be done on the whole image or on the extracted

patches. Among the possible methods for normalizing complete images, the following two are considered:

Minimum/maximum spreading: for minimum/maximum spreading, in a given image the brightest and the darkest pixel in the image are determined, i.e. the pixels with highest and lowest value. Let these values be v_{max} and v_{min} . Unless $v_{max} = v_{min}$, for each pixel p its value v is changed to $\frac{v - v_{min}}{v_{max} - v_{min}}$. Minimum/maximum spreading ensures that the full range of brightness values is used in the image. Nonetheless, as one can see in Figure 3.5, this does not lead to the desired result. The images remain almost unaffected, since already without normalization it contains very bright and dark areas.



Figure 3.4: Left: bright face. Right: dark face.

Histogram normalization: Applying histogram normalization, for a given image, a histogram $H = (h_1, \dots, h_n)$ with n bins is calculated, where n is equal or less than the number of possible values per pixel (typically 256 for gray-value images). The bins represent n discrete brightness values b_1, \dots, b_n . After inserting the pixels into H according to their brightness values and normalizing H , a cumulative histogram $\hat{H} = (\hat{h}_1, \dots, \hat{h}_n)$ is derived from H , where a bin \hat{h}_i is defined as $\hat{h}_i = \sum_{j \leq i} h_j$. Finally, if in the original image a pixel has a brightness value between b_{i-1} and b_i , its value is replaced by the relative count of the bin \hat{h}_i . Therefore, histogram normalization ensures that the brightness values are almost uniformly distributed. The problem with both approaches for brightness normalization presented above is that they normalize the whole image without focussing on the features of interest. Therefore, the brightness difference of these

features may be retained. We assume that the recognition performance is best if the features are normalized rather than the whole images. Thus, instead of the whole images, the patches have to be normalized. Still, the two previously mentioned methods are not likely to work.



Figure 3.5: Left: bright face after minimum/maximum spreading. Right: dark face after minimum/maximum spreading.



Figure 3.6: Left: bright face after histogram normalization. Right: dark face after histogram normalization.

Because the patches are small compared to the whole image, some of them are likely to cover homogeneous regions such that all their pixels have a similar brightness. In [Deselaers & Keysers+ 05b], normalizing patches is done by discarding the first PCA component of the patches after the PCA transformation, which approximates the effect of a brightness normalization as the “energy” of a patch, i.e. its overall brightness, is mainly in the first PCA coefficient. In the domain of face recognition, this approach has been taken, too [Martinez & Kak 01]. The effect is shown in Figure 3.7: from both the bright and the dark face shown before a patch depicting a part of the forehead is extracted, which is shown in (a) and (b). The difference in the brightness is obviously maintained by the PCA transformation, as shown in (c) and (d), where the patches are shown after PCA transformation and back-transformation of the first 40 PCA coefficients. Figure 3.7 (e) and (f) shows the back-transformed patches where the first of the 40 PCA coefficients has been discarded.

Apparently, the resulting patches have a similar brightness. Thus, unlike the two approaches presented first this one seems promising when it comes to normalizing the brightness of features in complex images.

3.3.3 Derivatives

Derivatives have been successfully used in feature recognition, for example in [Keysers 06]. Two types of derivatives are considered in this work: “Spatial” derivatives and derivatives using Sobel matrices. With spatial derivatives, the difference between the values of neighboring patches is taken into account. For each patch x of (quadratic) size w extracted at any extraction point (h_x, v_x) , two more patches x_{top} and x_{left} of the same size are extracted above and left to x , such that x and x_{top} resp. x and x_{left} are immediate neighbors, i.e. they have neither space in between nor do they overlap: x_{top} is located at $(h_x, v_x - w)$ and x_{left} is located at $(h_x - w, v_x)$. The feature vector of x is enlarged by the feature vector differences of x and x_{top} resp. x and x_{left} . Formally, if (x_1, \dots, x_n) is the feature vector of x without derivatives, and $(x_{top1}, \dots, x_{topn})$ resp. $(x_{left1}, \dots, x_{leftn})$ are the those for x_{top} resp. x_{left} , then the feature vector including the spatial derivative for x is defined as $(x_1, \dots, x_n, x_1 - x_{top1}, \dots, x_n - x_{topn}, x_1 - x_{left1}, \dots, x_n - x_{leftn})$. Thus, it is three times larger than the feature vector without spatial derivative. The feature vectors of the patches x_{top} and x_{left} are not furtherly considered, i.e. they are only determined to calculate the spatial derivative of x . Sobel matrices provide a mean to calculate local derivatives on the pixel level in images.




Figure 3.7: (a), (b): Two untransformed forehead patches with brightness difference. (c), (d): Forehead patches after PCA transformation and back-transformation with PCA coefficients 1 to 40: brightness difference remains. (e), (f): Forehead patches after PCA transformation and back-transformation with PCA coefficients 2 to 40: almost equal brightness.

Horizontal Sobel matrix	Vertical Sobel matrix
$\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$

Figure 3.8: 3×3 Sobel matrices.

Figure 3.8 shows 3×3 Sobel matrices for the horizontal and vertical direction. When convolving an image X with one of these matrices, horizontal resp. vertical edges and borders are emphasized. This effect can be

seen in Figure 3.9, which shows an airplane image in original, convolved with the horizontal 3×3 Sobel matrix, and convolved with the vertical 3×3 Sobel matrix.

EXPERIMENTAL RESULTS AND DISCUSSION

As discussed earlier while in the process of synthesizing or mapping a new High Resolution image from low resolution image. Patches are selected based on feature selection. Multi view features are exploited to embedded the high resolution patches to low resolution patches. Below figure displays the multi view features of the image at different levels of gradients.

Notation

The following table gives an overview of the symbols which are repeatedly used in this work

Table D.1: Symbols used in this work.

Symbol	Description
X	an image
x_l	feature vector representing the l -th patch in image X
$x_{n,l}$	feature vector representing the l -th patch in image X_n
L	number of patches extracted from each image
N	number of training images
$\{x_l^L\}$	set of L feature vectors $\{x_1, \dots, x_L\}$
y_l	position of the l -th patch in image X
$z_{\lambda,l}$	relative position of the l -th patch to the λ -th patch in image X
$\{z_l^L\}_l$	relative positions of the l -th patch to all other patches of the same image
r	decision rule
c	a cluster of feature vectors
μ_c, Σ_c	mean vector and covariance matrix of cluster c
C_k	number of clusters for class k
c_l	cluster whose center is most similar to the feature vector x_l
k	object class
k_n	object class of image X_n
f_i	feature function in log-linear models
$\lambda_{k,i}$	weight for feature i of class k in log-linear models
H	a histogram
d	number of dimensions in a histogram
v	number of different values per dimension in a histogram
M	number of bins in a histogram
c_i, \hat{c}_i	absolute and relative count for the i -th bin in a histogram
q	membership function for histograms

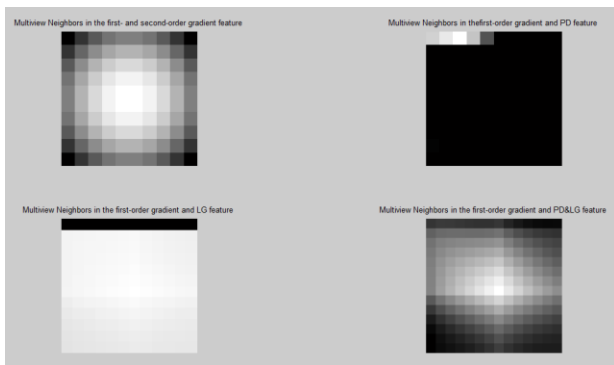


Fig 5.1 multiview features of the image at different levels of gradients.

Here each multi view feature is displayed based on intensity of the patch i.e. each patch represents the intrinsic value of patch.

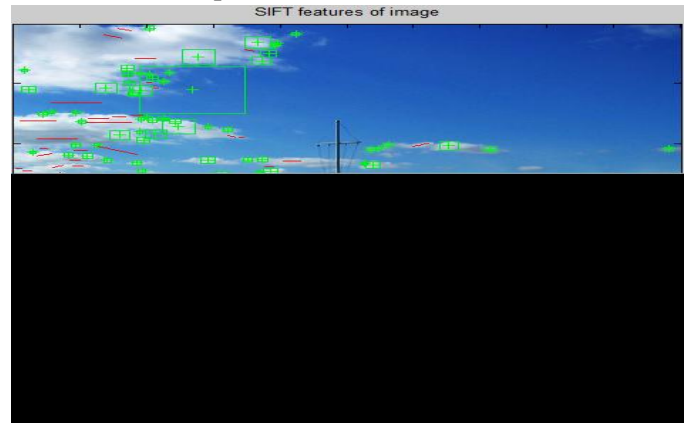


Fig.5.2 Feature point of the image based the intensity including colour value.

Above figure represents the feature points of the image based the intensity including colour value. Above figure consists of linear and closed feature points. These features are used to describe the variation in pixel a local window, and the variation along the cross direction respectively The PD feature can distinguish smooth patches from patches with textures and edges, and the LG feature can capture the detailed information in the horizontal and vertical directions.




Fig.5.3 Result of feature based patch image denoising model.

Above figure represents the results for feature based patch denoising model. From the above figure, it clearly describes that proposed scheme effectively denoising the image.

REFERENCES

- [1]. S. Agarwal, D. Roth. Learning a Sparse Representation for Feature Detection. Proc. ECCV, number 2353 in LNCS, pp. 113–127, Copenhagen, Denmark, May, June 2002.
- [2]. I. Assent, A. Wenning, T. Seidl. Approximation Techniques for Indexing the Earth Mover's Distance in Multimedia Databases. Proc. International Conference on Data Engineering, in press, 2006.
- [3]. A. L. Berger, S. D. Pietra, V. J. D. Pietra. A Maximum Entropy Approach to Natural Language Processing. Computational Linguistics, Vol. 22, No. 1, pp. 39–71, 1996.
- [4]. P. Bradley. Mathematical Programming Approaches to Machine Learning and Data Mining. Ph.D. thesis, University of Wisconsin, Computer Sciences Department, Madison, WI, USA, 1998. TR-98-11.
- [5]. J. Dahmen. Invariant Image Feature Recognition using Gaussian Mixture Densities. Ph.D. thesis, RWTH Aachen University, Aachen, Germany, Oct. 2001.
- [6]. T. Deselaers, D. Keysers, R. Paredes, E. Vidal, H. Ney. Local Representations for Multi-Feature Recognition. Proc. DAGM 2003, Pattern Recognition, 25th DAGM Symposium, number 2781 in Lecture Notes in Computer Science, pp. 305–312, Magdeburg, Germany, Sept. 2003. Springer Verlag.
- [7]. T. Deselaers, D. Keysers, H. Ney. Features for Image Retrieval – A Quantitative Comparison. Proc. DAGM 2004, Pattern Recognition, 26th DAGM Symposium, number 3175 in Lecture Notes in Computer Science, pp. 228–236, Tübingen, Germany, Sept. 2004.
- [8]. T. Deselaers, D. Keysers, H. Ney. Discriminative Training for Feature Recognition using Image Patches. Proc. IEEE Conference on Computer Vision and Pattern Recognition, Vol. 2, pp. 157–162, San Diego, CA, June 2005.
- [9]. T. Deselaers, D. Keysers, H. Ney. Improving a Discriminative Approach to Feature Recognition using Image Patches. Proc. DAGM 2005, Pattern Recognition, 26th DAGM Symposium, number 3663 in Lecture Notes in Computer Science, pp. 326–333, Vienna, Austria, Aug. 2005.
- [10]. T. Deselaers. Features for Image Retrieval. Diploma thesis, Lehrstuhl für Informatik VI, RWTH Aachen University, Aachen, Germany, Dec. 2003.
- [11]. G. Dorko, C. Schmid. Selection of Scale-Invariant Parts for Feature Class Recognition. Proc. International Conference on Computer Vision, Vol. 1, pp. 634–640, Nice, France, Oct. 2003.